Red Hat Linux 6.1

The Official Red Hat Linux Getting Started Guide

Contents

Red Hat Linux 6.1

Chapter 7   Q & A: Quick Answers to Common Questions

Part II The GNOME User's Guide

Chapter 8   The GNOME User's Guide

Chapter 9   GNOME Quick Start

Chapter 10   Window Managers and GNOME

Chapter 11   The Enlightenment Window Manager

Introduction

Welcome to the Official Red Hat Linux Getting Started Guide!

At Red Hat, we believe we offer the best Linux distribution on the market. We hope you'll agree that the time and the money you spent to install and work with Red Hat Linux was well spent, indeed.

If you count yourself among the many who are discovering Red Hat Linux for the first time, this book is for you!

Inside, you'll find valuable tips which will help you get acquainted with GNOME, your new default desktop environment. You'll also learn the fundamentals of navigating and managing files and directories from the **shell prompt**, or console. Among the new features of Red Hat Linux 6.1 you'll find the RH PPP Dialer, also called RP3; inside, you'll learn how to use this application to get online quickly and easily.

In short this guide picks up where the Official Red Hat Linux Installation Guide leaves off: helping you understand a bit about your new operating system.

But first, a little background…

What is Linux?

Back in August of 1991, a student from Finland began a post to the comp.os.minix newsgroup with the words:

```
Hello everybody out there using minix -

I'm doing a (free) operating system (just a hobby, won't
be big and professional like gnu) for 386(486) AT clones.
```

The student was Linus Torvalds, and the "hobby" he spoke of eventually became what we know today as Linux.

A full-featured POSIX-like operating system, Linux has been developed not just by Linus, but by hundreds of programmers around the world.

This massive, world-wide development effort is largely uncoordinated. Sure, Linus calls the shots where the kernel is concerned, but Linux is more than just the kernel. There's no management infrastructure; a student in Russia gets a new motherboard, and writes a driver to support a neat feature the motherboard has. A system administrator in Maryland needs backup software, writes it, and gives it away to anyone that needs it. The right things just seem to happen at the right time.

And Linux can be obtained for absolutely no money. That's right: Most of the software is available (at no charge) to anyone with the time and inclination to download it. But not everyone has that much time…

What is Red Hat Linux?

Enter a group of programmers based in North Carolina. Their goal was to make it easier for people to give Linux a try. Like many other such groups, their approach was to bundle all the necessary bits and pieces into a cohesive **distribution**, relieving "newbies" from some of the more esoteric aspects of bootstrapping a new operating system on their PCs.

However, unlike other distributions, this one was fundamentally different. The difference? Instead of being a snapshot of a hard disk that had a working copy of Linux on it, or a set of diskettes from which different parts of the operating system could be dumped, this distribution was based on **packages**.

Software development in the Linux world is fast-paced, so new versions of old software come out continually. With other distributions, upgrading software was painful -- a complete upgrade usually meant deleting everything on your hard drive and starting over.

Each package provided a different piece of software, fully tested, configured, and ready to run. Want to try a new editor? Download the package and install it. In seconds, you can give it a try. Don't like it? Issue a single command, and the package is removed.

If that was all there was to it, this distribution would be pretty nifty. But being package-based meant there was one additional advantage:

This Linux distribution could be easily upgraded.

By now you've probably guessed that the group of programmers in North Carolina is Red Hat, and the package-based distribution is Red Hat Linux.

Since Red Hat Linux's introduction in the summer of 1994, Linux and Red Hat have grown by leaps and bounds. Much has changed; support for more esoteric hardware, huge increases in reliability, and the growing use of Linux by companies around the world.

But much still remains the same. Linux is still developed by people world-wide; Linus is still involved. Red Hat is still headquartered in North Carolina; still trying to make Linux easier for people to use.

And Red Hat Linux is still package-based; always has been, always will be.

Since the release of version 4.0, Red Hat Linux runs on three leading computing platforms: Intel compatible PCs, Digital Alpha computers, and Sun SPARC equipment. Our unified source tree and the benefits of RPM (Red Hat Package Management) technology enable us to deploy Red Hat Linux for each platform with a minimum of effort. This, in turn, enables our users to manage and port software between these platforms as easily as possible.

We make Red Hat Linux available by unrestricted FTP from our site and many mirror sites on the Internet. Red Hat Linux is also available on CD-ROM. For current information on our product offerings and links to other Linux resources please check Red Hat's website at http://www.redhat.com.

On most systems, Red Hat Linux is easy to install; the installation program can walk you through the process in as little as 15 minutes. The system itself is very flexible. With RPM, you can install and uninstall individual software packages with minimal effort.

Because of RPM, Red Hat Linux is also easy to maintain -- package installations can be verified and corrected, and packages can be installed and uninstalled simply and reliably. Furthermore, Red Hat Linux is easy to administer. Included are a rich set of administrative tools which reduce the hassle of everyday system administration. Complete source code is provided for the freely distributable components of the system.

About This Book

This book is divided into two parts:

- The Basics

- The *GNOME User's Guide*.

Let's look at each part in a bit more detail.

The Basics

Here, you'll find an introduction to some of the more important features of your Red Hat Linux system, as well as some of the immediate tasks you must complete.

- Quick View of GNOME -- A brief look at GNOME, the default desktop environment for your Red Hat Linux system.

- File Essentials -- Fundamental information you should have to navigate and manage files and directories on your Red Hat Linux system. If you didn't take the opportunity to create one or more user accounts during your installation of Red Hat Linux, you can learn how in these chapters. You'll also learn how to modify accounts that you've already made.

- Getting Online -- Creating a PPP connection has never been easier with Red Hat Linux. Here's where you might find all the information you'll need to create and modify PPP connections.

- Using RPM -- RPM technology means convenience, power and safety for you. Here's where you can gain an understanding of how to get results with RPM, in both the GUI and at the shell prompt.

The *GNOME User's Guide*

An indispensable resource for navigating and customizing GNOME, the *GNOME User's Guide* begins with Chapter 8, *The GNOME User's Guide* and continues through Chapter 21, *The GNOME Calendar*.

GNOME stands for *G*NU *N*etwork *O*bject *M*odel *E*nvironment. That's a fancy acronym, but it translates into a pleasing environment which offers all the power of Linux. GNOME is the default X Window System environment for Red Hat Linux 6.1.

In the *GNOME User's Guide*, you'll find ways to create, move and copy files, investigate your new system and more -- all within a pleasing graphical environment.

Here's a preview of what you'll find in these chapters:

- An Introduction to GNOME -- Learn a little about the history and ideas which have helped create GNOME.

- GNOME Quick Start -- Get a "quick start" on familiarizing yourself with the GNOME desktop.

- Window Managers and GNOME -- Learn a little about window managers, and which window managers you can use that work well with GNOME.

- Using the GNOME Panel -- Learn the ins and outs of GNOME Panel, from which you can find and start your system's applications. Highly flexible, the Panel allows you to add favorite applications, change its location, add new panels and more.

- The GNOME Desktop -- Your GNOME desktop is more than just a pretty face. Here's where you can find out how to "drag and drop" applications, store files, create folders, switch desktops and more.

- The File Manager -- Here, you'll learn how to view and manipulate files -- both on your system and on the Internet.

- The Main Menu -- Even though it gives you easy access to applications, you're not stuck with the main menu's default configuration. You can learn how to make the menu work the way *you* work.

- The GNOME Control Center -- As advertised, the GNOME Control Center allows you to control the way your environment looks and "feels." You'll learn how to tweak your environment in many ways.

- Session Management -- Learn how you can automatically start favorite applications when you log in.

- Using Panel Applets -- System and network monitors, multimedia enhancements, games and more. Applets are small applications which can add productivity and fun to your environment. You'll learn all about how to work with applets in the following related chapters.

You'll also find quite a few translations of the *GNOME User's Guide*, as well as the latest GNOME documentation and software at the official website: http://www.gnome.org/.

Now, you've probably got a question or two…

Q & A and the Glossary

When you care only for answers, turn to these chapters. Many of the entries you'll find in Chapter 7, *Q & A: Quick Answers to Common Questions* concern answers to some of the most common new user questions, such as:

- How do you access a floppy disk?
- What happens if you forget your root password? (Don't panic! You can create a new one.)
- Can you choose which **services** you want when you start your Red Hat Linux system?
- How do you access files on a Windows partition?

In the Glossary, Appendix A, *Glossary*, you'll find quick definitions to plenty of words and phrases that can vex any new user.

More to Come

As Linux evolves, so does the support you'll find for Red Hat Linux. The Official Red Hat Linux Getting Started Guide is part of that support -- and evolution. In coming editions, expect to find more essential information to help you get the utmost from your Red Hat Linux system.

And here's how you can help.

Send in Your Feedback

If you'd like to make suggestions about the Official Red Hat Linux Getting Started Guide -- from words to add in the Glossary to typos or subjects you'd like to see covered -- please mention this guide's identifier:

GSG(EN)-3.0-Print-RHI (09/99)

That way we'll know exactly which version of the guide you have. You can send mail to:

`docs@redhat.com`

Sign Up for Support

If you have an official edition of Red Hat Linux 6.1, please remember to sign up for the benefits you're entitled to as a Red Hat customer.

You'll be entitled to any or all of the following benefits, depending upon the Official Red Hat Linux product you purchased:

- Official Red Hat support -- Get help with your installation questions from Red Hat, Inc.'s support team.

- Priority FTP access -- No more late-night visits to congested mirror sites. Owners of Red Hat Linux 6.1 receive free access to priority.redhat.com, Red Hat's preferred customer FTP service, offering high bandwidth connections day and night.

- Red Hat Update Agent -- Receive e-mail directly from Red Hat as soon as updated RPMs are available. Use Update Agent filters to receive notification and quickly download updated packages about those subjects that interest you. Also receive -- automatically -- kernel updates, security updates and other packages.

- Under the Brim: The Official Red Hat E-Newsletter -- Every month, get the latest news and product information directly from Red Hat.

To sign up, go to http://www.redhat.com/now. You'll find your **Personal Product ID** on a red and white card in your Official Red Hat Linux box.

To sign up, go to www.redhat.com/now. You'll find your **Personal Product ID** on the red and white card in your Official Red Hat Linux box. To read more about technical support for Official Red Hat Linux, refer to the Appendix in the Official Red Hat Linux Installation Guide.

Good luck, and enjoy your new Red Hat Linux system!

The Red Hat Documentation Team

Part I The Basics

1 GNOME in 60 Seconds (or a little longer)

GNOME is the default desktop environment in Red Hat Linux 6.1. Flexible and user-friendly, GNOME provides both the power of your new Red Hat Linux system and a comfortable GUI.

In this section, we'll briefly cover some essentials of your new environment to help you feel right at home. For greater detail about GNOME, please turn to Chapter 8, *The GNOME User's Guide* later in this book; the guide is also accessible from your desktop in the GNOME Help Browser or from the Web at http://www.gnome.org.

1.1 An Overview

When you initially begin your **X Window session**, among the first items you'll see are:

- **The Desktop** -- The main workspace of your environment. Windows are placed on the Desktop; **launchers** are also found here for applications, documents or other data. Double-clicking will open a folder or start an application.

- **Folders** -- Containing links to documentation such as support, help and reference guides, either on your installed system (locally) or on the Web. Folders (as well as launchers to other applications) can be added and deleted from the Desktop. See Figure 1–2, *Folders and Launchers on the Desktop* for an example of folders and launchers on the desktop.

- **The Panel** -- The long bar across the bottom of the Desktop serves as a quick launcher for **applets**; the stylized **GNOME footprint** (called the **Main Menu Button**) leads to the menu of applications. Single-clicking on the **Main Menu Button** will display the included applications on your system.

1.1.1 Using the Mouse

A three-button mouse offers the greatest flexibility in GNOME. If you have a two-button mouse, however, and you selected three-button **emulation**, simply hold down both left and right mouse buttons at the same time.

Pressing the middle button of your mouse on the desktop will reveal a GNOME applications menu, help menus and configuration menus for Enlightenment, the default GNOME **window manager**.

Using the left mouse button selects an item to open or drag to another location; by holding down the left mouse button while you drag the cursor over items, you can select multiple items on the desktop.

Using the right mouse button reveals a menu of actions which you can take. Right-clicking on the desktop shows you options such as adding desktop items, rescanning the desktop and more. Right-clicking on a folder or application reveals a menu related to actions you can take with your selection, such as deletion, properties and more.

1.1.2 Drag and Drop

You can **drag and drop** items in GNOME, meaning that you can "carry" data or applications from one area of the desktop to another, "drop" items in folders or on applications, move or copy information and applications.

1.2 The Desktop

Your first view of GNOME will look something like Figure 1–1, *A GNOME Desktop*.

Figure 1–1   A GNOME Desktop

Double-click on any of the following items on your desktop to start or open the associated link:

- Launchers -- Icons which lead to websites (such as the Red Hat, Inc. website at http://www.redhat.com/), support documentation on your installed system and more.

- Folders -- Containing contents of directories, launchers to documentation or applications. The folder **Home directory**, for example, leads to your specific directory in /home on the system (for example, if your login directory was bob, the **Home directory** folder on the desktop would lead to your login directory, /home/bob).

You may also see a warning dialog urging caution, if you logged in as **root**, because you hadn't created your user account during your installation of Red Hat Linux. Click on the **OK** button to dismiss this warning dialog (take the warning to heart, however…).

Figure 1–2   Folders and Launchers on the Desktop

Double-click on a launcher to start the associated file (see Figure 1–2, *Folders and Launchers on the Desktop*). If it's an application, it will start; if it's a directory, it will open and if it's data, that data will be displayed in its associated program (a spreadsheet, for example, will open in Gnumeric).

1.3 Folders

Folders on the desktop can store applications, data such as text files or spreadsheets and other items.

Double-clicking on a folder will open it and reveal its contents (see Figure 1–3, *Contents of an Open Folder*).

Figure 1–3   Contents of an Open Folder
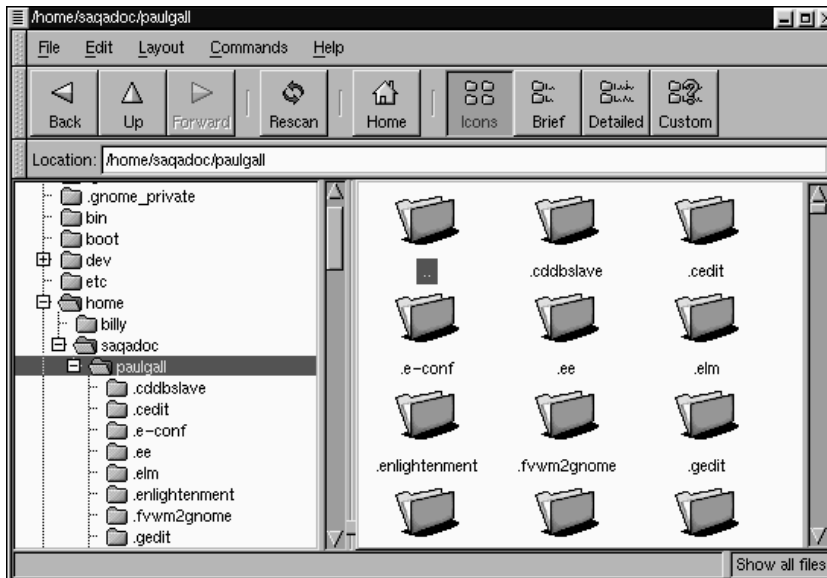
In fact, folders are displayed in the GNOME File Manager, one of the more powerful tools at your disposal -- both as a user and as a system administrator. You can learn more about the File Manager in the *GNOME User's Guide* later in this manual.

1.4 The Panel

Like much of GNOME, the Panel is highly configurable (see Figure 1–4, *The Panel at the Bottom of the Desktop* for an example panel).

From the **Panel**, you can launch applications and **applets**, small applications and utilities which are meant to be docked into the panel. Among the applets already running on the **Panel** are the Pager and a clock applet. Other applets which can enhance productivity and enjoyment range from a mail checker and a CD player to system monitors.

Figure 1–4   The Panel at the Bottom of the Desktop

- To add an applet -- Right-click on the **Panel**, as in Figure 1–5, *Adding an Applet*. Select one of the applets listed in the sub-menu, and it will be docked into the **Panel**.

- Changing the **Panel**-- You can add more than one **Panel** to your desktop, or convert the current orientation of the **Panel** from horizontal to vertical, or "shrink" the **Panel** to appear in a corner of the Desktop, rather than extend across the screen.

Figure 1–5   Adding an Applet
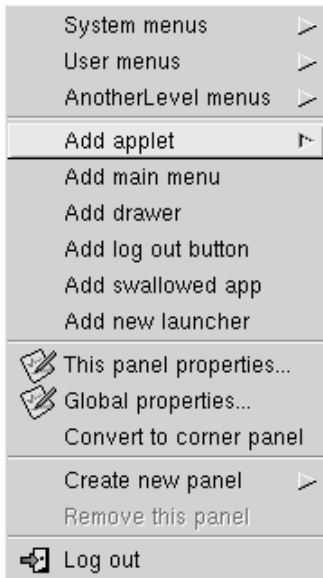
Applications, folders and applets can be anchored on the **Panel**, so that they can be launched with a single click of the button.

1.5 Finding Help

One of the most useful utilities for the GNOME user can be GNOME Help Browser (see
Figure 1–6, *The GNOME Help Browser*).

In fact, since this application displays not only GNOME-related help but also **info
pages** and **man pages**, the GNOME Help Browser can be an invaluable resource for learn-
ing about both GNOME and your Red Hat Linux system. (Info pages and man pages
are concise help documents which are usually installed on a Linux user's system. To
learn more about these documents, read on to Chapter 2, *Welcome to Linux*.)

Figure 1–6   The GNOME Help Browser

You can open the GNOME Help Browser by double-clicking on its launcher on the desktop
or from the **Panel** by dragging your cursor to the item **Help System** from the **Main Menu
Button**.

You can learn more about the Help Browser in the GNOME User's Guide later in this book.

Now, turn to Chapter 2, *Welcome to Linux*, to learn more about your new Red Hat Linux system, and which important first steps you should take.

2 Welcome to Linux

What began in 1991 as a hobby for a young Finnish student named Linus Torvalds has ballooned from a "hacker's darling" into an important tool for both home and business users.

Just six years ago, there were an estimated 100,000 users. Today, millions of users worldwide depend on Linux to manage finances, use and control Internet services, create artwork and more. That number is rapidly growing; every day, new users are discovering the power and potential of Linux.

This free operating system is a multitasking, multi-user environment that has superior memory management, great security features, and more.

In other words, power and, once you become more comfortable with Linux, ease of use.

Tip

Linux is most frequently pronounced with the short "i" and the accent on the first syllable, as in "LIH-nucks."

What do you do next? Relax.

In the chapters that follow, we'll show you the basics of your new operating system, from fundamental tasks such as creating a new **user account** and getting online to tips on how to install and manage new software.

2.1 A Note About Conventions

At the time you installed your Red Hat Linux system, you were given the option of working entirely in a **GUI,** or graphical user environment, such as GNOME, or logging in from **console mode,**, which is non-graphical.

If you're like many new Linux users, you're familiar with graphical environments such as Microsoft Windows, Apple Macintosh or IBM's OS/2.

2.1.1 How to Find an Xterm

You likely chose to work in a GUI when you installed Red Hat Linux 6.1.

You'll find plenty of opportunities to "point and click" on applications, either on your desktop or from the **Panel**, at the bottom of your desktop. But we're also going to try to dedicate some time to working from the **shell prompt**.

Why? Because at the same time you accomplish tasks, you can learn a little more about how your Red Hat Linux system works.

Tip

Unlike a graphical presentation, a "shell prompt" is the way you can type commands directly to the "shell." You need a shell for Linux, because it's the tool you use to interact with your operating system. You'll find more information about your shell in Chapter 4, *Managing Files and Directories*.

Figure 2–1   The Main Menu Button

There are plenty of ways to get a shell prompt, depending on the kind of graphical environment you're using, such as GNOME. Often, just by right- or left-clicking in a blank space on your desktop, you'll see a reference to an **Xterm**, or **terminal**. If you right-click on an empty space on your GNOME desktop, a menu will direct you to **New** => **Terminal**. Allow your cursor to highlight the **Terminal** entry, then release the button.

By "dragging" your cursor over that item and releasing the button you will open a shell prompt window. Other times, you'll find you can get to a shell prompt through the menu on your desktop.

In addition to **Xterm**, other references which will give you the shell prompt include:

- terminal emulator window

- GNOME terminal

- Color Xterm

Since we'll use GNOME as our example, you'll find an alternative method of opening an Xterm window is from the **Main Menu.**

You can do this by taking your cursor to the **Main Menu Button,** the stylized GNOME "footprint" on the Panel.

Now, left-click once on the **Main Menu Button** (see Figure 2–1, *The Main Menu Button*), and a menu of "folders" will pop up. These folders represent categories of various software groups on your system. There are utilities, graphics programs, Internet applications and many more.

Once the menu items appear, "drag" the cursor to the **Utilities** section of the menu by holding down the mouse button while raising the mouse to the Utilities folder.

When you're over the folder, a new menu pops up to its right, revealing a submenu of applications in the Utilities category.

Here, in the first entries of this new submenu, there is a choice of terminal windows: **Regular Xterm**, **Color Xterm** and **GNOME terminal** (as shown in Figure 2–2, *Shell Prompt Menus in Utilities*). To get a shell prompt, position the cursor over the terminal window of your choice, and release the mouse button.

Figure 2–2   Shell Prompt Menus in Utilities
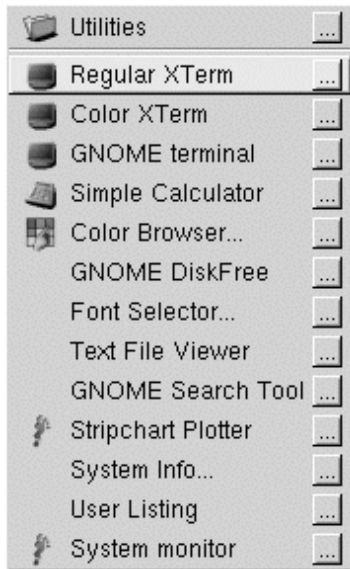
Tip

Of the choices in the **Panel** => **Utilities** menu, the GNOME terminal offers the ability to easily configure background color and other preferences.

Now, it's time to take your first steps.

2.2 The Root of the Matter

When your Red Hat Linux system starts, you'll see an array of messages speeding past you on the screen. Many of these messages simply tell you what services are starting on your computer.

Tip

Want to read those startup messages more closely? At a shell prompt, type `dmesg | more`. You'll be able to read the file one screen at a time. To move forward, press the [Spacebar]; to quit, press [Q].

Finally, you'll come to "the login prompt" (as shown in Figure 2–3, *Sample Screen of the Graphical Login Prompt*).

You'll find:

```
Login:
Password:
```

At this point, some new users can easily feel rattled, but don't panic. Instead, think back: When you installed Red Hat Linux 6.1, you were asked for a **root password**.

Figure 2–3   Sample Screen of the Graphical Login Prompt
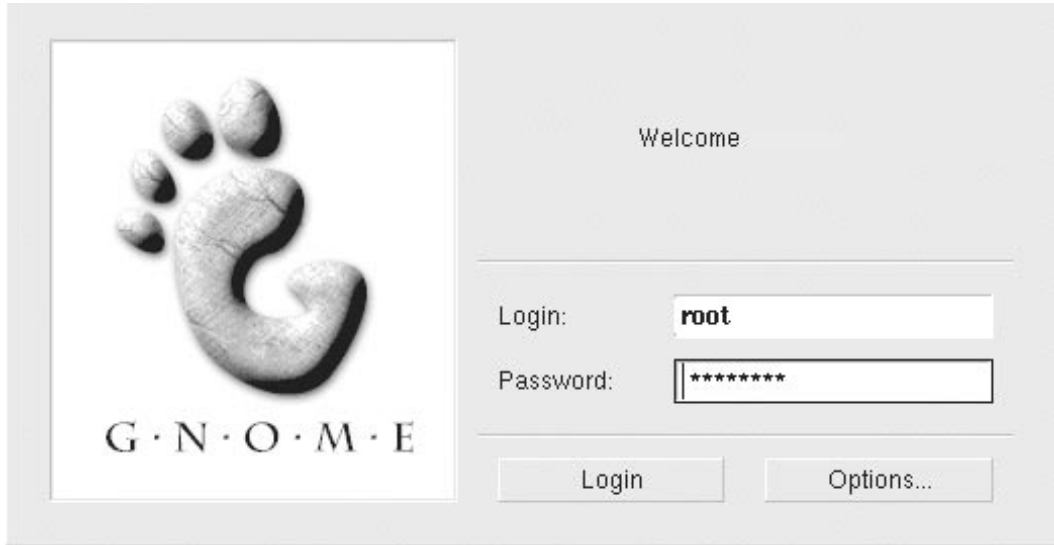
In detail, that is the password you were asked to choose to log in to your root account.

During your installation of Red Hat Linux 6.1, you also had the opportunity to create a **user account**. If you didn't create your user account at that time, there's nothing to worry about: you'll see how to create a user account in just a few minutes.

Your user account is the account you'll log into for just about everything, except system maintenance and other administrative tasks. For these system tasks, you use the root account.

When you log in -- either in the root account or other accounts -- you're introducing yourself to the system. The root account, unlike all other accounts for your system, has access to everything. Also known as the **superuser**, the root can control everything the system does.

For now, let's log in as root, take a brief look around, create a user account and then log out.

Go ahead and login; at the `Login:` prompt just type:

> **root**

and press [Enter] or the [Tab] key.

Tip

Case matters. Linux, like UNIX, makes a distinction be-
tween uppercase and lowercase letters. So *root* is not the
same as *Root*. In fact, as far as Linux is concerned, they're
two different accounts.

Don't worry about mistakes when you log in; you can always use the [Delete] key to
start over.

When you're asked, type in the password you chose when you installed Red Hat
Linux. You won't see your password on the screen as you type; that's just one of the
security features of your Linux system.

Tip

Be sure to type commands exactly as you see them -- spaces,
dashes and all. To Linux, an extra space or letter can make
all the difference in the world.

When you're finished typing in your password, press [Enter].

You'll be presented with your new desktop (similar to Figure 2–4, *Sample GNOME Desktop*). Once you become more comfortable with your new operating system and with GNOME, you'll probably be able to fill up that desktop quickly with applications.

Figure 2–4   Sample GNOME Desktop

There's plenty of space to fill up, too; not just on your current desktop but on numerous desktops -- four by default.

Tip

To learn specifics about GNOME, including applica-
tions and navigation, turn to the *GNOME User's Guide*
section of this manual. Or, for the latest GNOME doc-
umentation and applications, visit the GNOME website,
http://www.gnome.org.

For now, let's concentrate on your current desktop.

Go ahead and look around. You can begin by double-clicking with the left mouse
button on the file folder called **Home directory** on the desktop.

Here, you'll find icons representing the various directories and files on your system.

From the panel on the bottom, left-click once with your mouse, and you can begin
to investigate some of the applications which have been included with your environ-
ment.

From here, you can find ways to customize your workspace, search for files, write
letters or other documents, start spreadsheets and more.

But before you get too daring…

While you're logged into the root account, avoid the temptation to make any changes
to files or directories unless you know exactly what you're doing!

Here's why: Whenever the system recognizes you as the root account you're allowed
to do just about anything -- change configuration files, make new directories, create
and manage accounts for users who are allowed to use your computer and more.

That kind of power comes with a price, and tinkering around with configuration files,
accounts and directories can easily lead to disaster.

So how are you supposed to operate safely?

That's where the user account come in handy. With a user account, you can work and
play with the assurance that you're not damaging your system.

2.3 How to Quit

When you're finished looking around for the first time, you can log out to quit your session (see Figure 2–5, *Locating the* **Log out** *Selection*).

Figure 2–5   Locating the Log out Selection
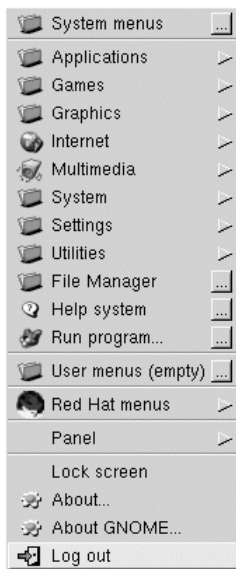
When you log out, you'll be returned to the opening screen you found when your system started.

To log out, just go to the **Main Menu Button** on the **Panel,** left-click once and select **Log out**.

You'll be presented with a box, asking you whether you want to log out. Select **Yes**.

After a few moments, you'll be returned to the log in screen.

Summary

At **Login** - type `root`. At **Password** - type `yourroot-password`. To quit - Left-click on the **Main Menu Button** => **Log out** => **Yes**.

2.4 From Console Mode

When you were installing Red Hat Linux 6.1, you were given the option of starting from a graphical or console (that is, non-graphical) screen.

If you chose not to automatically start your computer in a graphical environment, you'll find a somewhat daunting, almost blank screen which will show you something like

```
Red Hat Linux release 6.1
Kernel 2.2.x on an i686
login:
```

You can log in by typing **root** at the Login: prompt. Then, when Password: appears, type the password you chose at the time you installed Red Hat Linux 6.1.

Tip

Just like the graphical login screen, don't expect to see your password "echoed" when you type in your password. Making sure that your password isn't seen is just one of your new system's many security features.

Now, you'll find a single shell prompt, which will appear similar to:

```
[root@localhost root]#
```

Figure 2–6   Sample Shell Prompt After Login
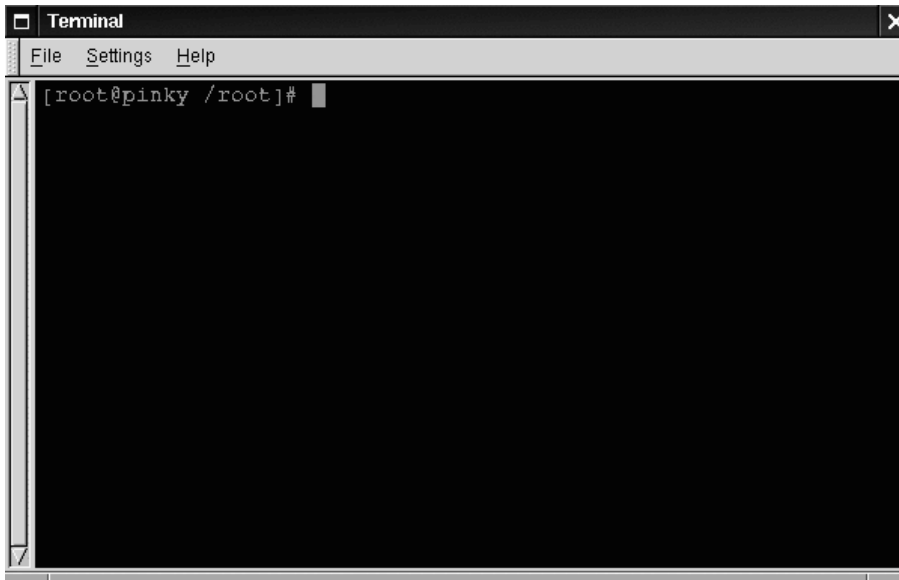
This tells you that you're logged in as root and in the directory called root (as shown in Figure 2–6, *Sample Shell Prompt After Login*).

You can exit at any time simply by typing logout or exit.

Tip

You can also press the [Ctrl]-[D] keys at the same time to return
you to the login prompt

2.5 X Marks the Spot

If you installed the X Window System (also known simply as X) at the time of your
Red Hat Linux installation, you've got a pleasing, graphical environment in which to
work.

If you didn't install X at that time, and you wish to use the X Window System, your
best bet is to return to the CD at this time and re-install Red Hat Linux 6.1. (Sigh…)

Certainly, there *are* other ways of installing X, but if you're fairly new to Red Hat
Linux, and if you're starting out with a brand new installation, you'll find it takes less
time -- and frustration -- to simply redo the installation.

If you did install the X Window System -- but didn't choose to start the X Window
System automatically, you're still ready to go.

At the shell prompt, type:

```
startx
```

and the X Window System will begin.

Please Note

You're logged into the system as the root account, also
known as superuser. There's a reason the root account is
known as superuser: In this account, you can make changes
to just about anything. Unless you know what you're doing,
you can easily harm your system by mistakenly changing
settings.

Although you may be tempted to modify files or directories, you should resist making
any changes until you've created a user account.

To log out of X, bring your mouse cursor to the **GNOME panel**, then left-click on the
Main Menu Button. A menu of applications, utilities, games and other programs will
pop up.

Drag your cursor to the item labeled **Log out** (as shown in Figure 2–5, *Locating the*
Log out *Selection*).

Tip

You can "drag" your cursor by keeping the mouse button
depressed with your finger while moving the cursor to your
selected item. Once the cursor is over the item on the menu,
releasing the mouse button will start the program.

Now, a separate window will appear, asking you to confirm your decision to log out.

Click on **Yes**, and you will return to your original shell prompt.

Whenever you want to start another X session, just type startx from the prompt.

If you'd rather log out of both your X session and your account, you can choose the options at the confirmation window (the window which asks **Really Log Out?** when you specify that your want to log out).

Specifically, in this window, you can:

- **Logout** -- End your X session and return you to the shell prompt. Or, if you're beginning with the graphical login screen, you'll return to that screen.

- **Halt** -- Shut down the system by stopping all the **services** in your system so you can turn off your machine.

- **Reboot** -- Shut down and restart your system.

When you select **Halt** or **Reboot**, you'll be asked to provide your password (whichever password you used to log in to the account). Just use [Enter] or click the **OK** button and your machine will shut down or reboot.

Regardless of the option you choose -- logout, halt or reboot -- you should select **Save Current Setup** in the logout window if you want your desktop environment to remain just as you left it in your current session.

Summary

At the prompt, type `startx`; to exit -- **GNOME panel** => **Log out**

2.6 A Brand New You

When you installed Red Hat Linux 6.1, you were given the option of creating a **user account**, in addition to specifying the root password. Such a handy feature can help you quickly get up and running in your own user account. If you didn't elect to create such an account, it's a pretty simple process.

Let's create a user account now. It can be your first user account or an additional user account. Once you're familiar with the process of adding user accounts, you can later edit those user accounts, to specify what the user can and can't have access to.

If you're familiar with MS-DOS or, to a lesser extent, Windows 98, you might be a little befuddled by the requirement of creating a user account.

After all, if you can navigate the system and use programs in your root account, you might think that having two accounts on a single machine is excessive.

Nothing could be further from the truth. Here's why:

Linux is a multi-tasking, multi-user system, which means it can safely and securely accommodate many users at one time, performing plenty of tasks each user requires. But only one account can be root -- capable of changing the way the operating system works.

Because "rooting around" can easily lead to havoc, it's important to safeguard against accidents. That's why just about every Linux user -- even if they're the system administrator -- has their own user account.

Once you're logged in as root, you have two ways to conveniently add a user to the system: from within X and from the shell prompt.

Both methods are quick and painless.

Let's say that the account you want to choose is called "billy."

From X:

One of the most powerful tools you can use for system administration is **Linuxconf**. You can use Linuxconf for adding and manipulating accounts, monitoring system activities and plenty of other system features.

Figure 2–7   Finding the Users Accounts Entry in Linuxconf
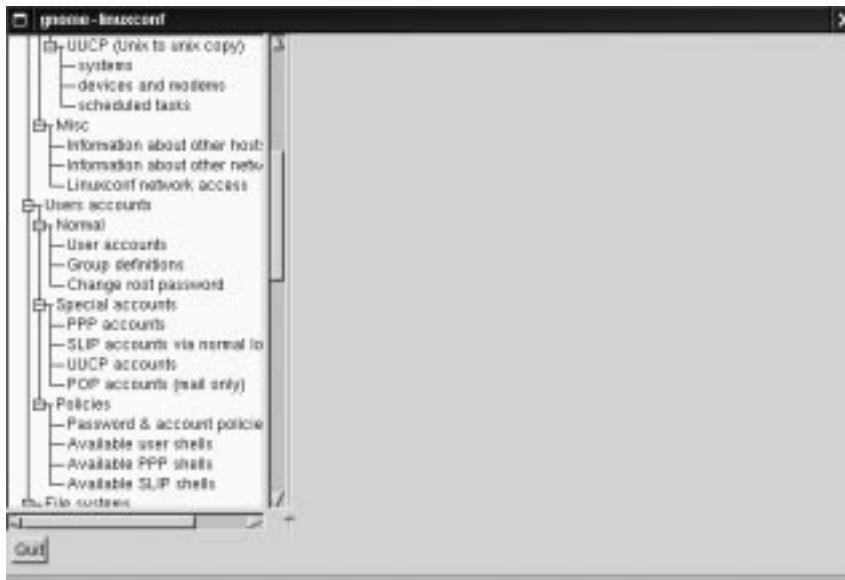
Completely documenting all the features of this utility would take more space than we have here. For a more detailed look at the application's features -- including greater depth on manipulating accounts -- turn to the *System Configuration with Linuxconf* chapter in the Official Red Hat Linux Reference Guide.

Tip

You can learn more about Linuxconf by visiting the utility's official website: http://www.solucorp.qc.ca/linuxconf/.

One of the easiest ways to access Linuxconf is from the shell prompt in an Xterm window. At the prompt, type:

```
linuxconf
```

We want to add an account, so scroll about a third of the way down the menu in the left panel, to the entry marked **Users accounts**. If the entry has a "**+**" next to it, go ahead and click on the "**+**".

Tip

Whenever you want to quit an Xterm window, all you've got to do is type **exit** at the prompt or, in some window managers, click on the **X** in the upper right corner of the window.

Now, the menu will expand to show entries in the **Users accounts** listing (as shown below).

The subentries will look like the following:

```
 |
-Users accounts
 |
 Normal
 |
 User accounts
 |
 Group definitions
 |
 Change root password
```

Left-click with your mouse button on the **User accounts** entry, under **Normal**.

In the right panel, you'll now see a box of the current user accounts (as shown in Figure 2–8, *Users Accounts in Linuxconf*).

Figure 2–8   Users Accounts in Linuxconf
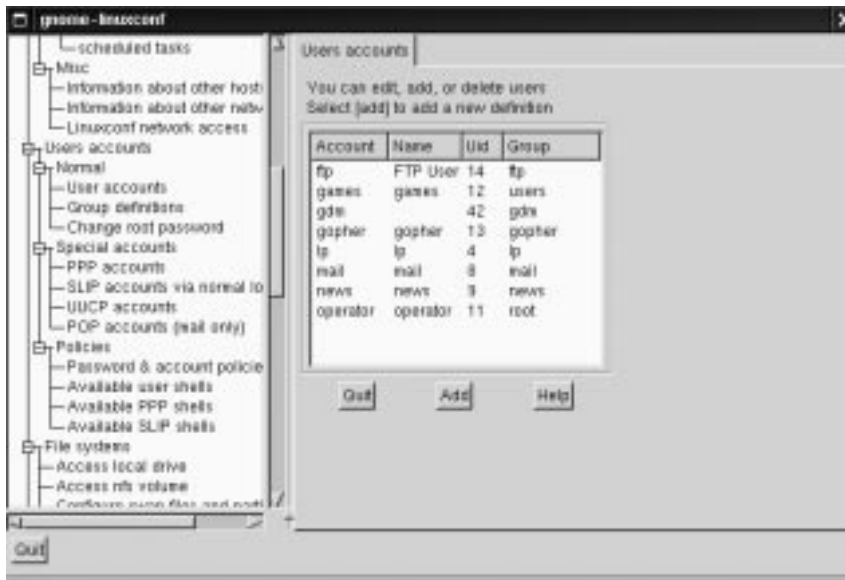
Toward the bottom of the right panel, click on the **Add** button, between the **Quit** and **Help** buttons.

Now, we'll see a dialog called **User account creation**. In here, we're going to fill in: **Login name**; **Full name** and **group**.

Make sure the button is indented next to the statement **The account is enabled**.

Now, let's type in a login name. It should be easy to remember (it's the password that should be complex, but more about that later …). Then, you can type in your full name.

Figure 2–9   Adding a User in Linuxconf
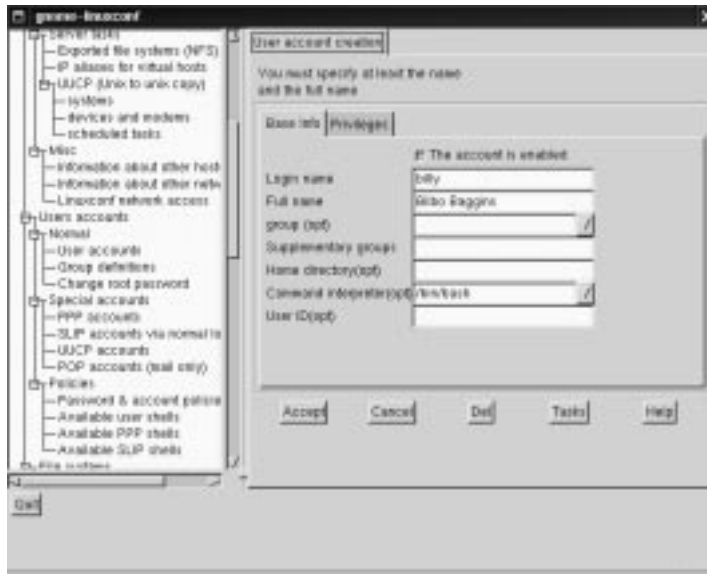
Tip

Everyone's account belongs to at least one group. Groups are used to determine file access permissions. Unless you specify a group, the default group for your user account will be the login name you choose (for example, a group called **billy**).

When you're finished, your entries should look like the following:

```
Login name.........billy

Full name..........Bilbo Baggins

group (opt)........billy
```

Now, just click on the button marked **Accept**.

We're almost finished. Next, we've got to come up with a password.

Passwords are one of the best methods to safeguard against prying eyes or malicious behavior. If you've got a secure password, which only you know, you've taken a big step in your system's security.

Figure 2–10 Creating a User Account Password in Linuxconf
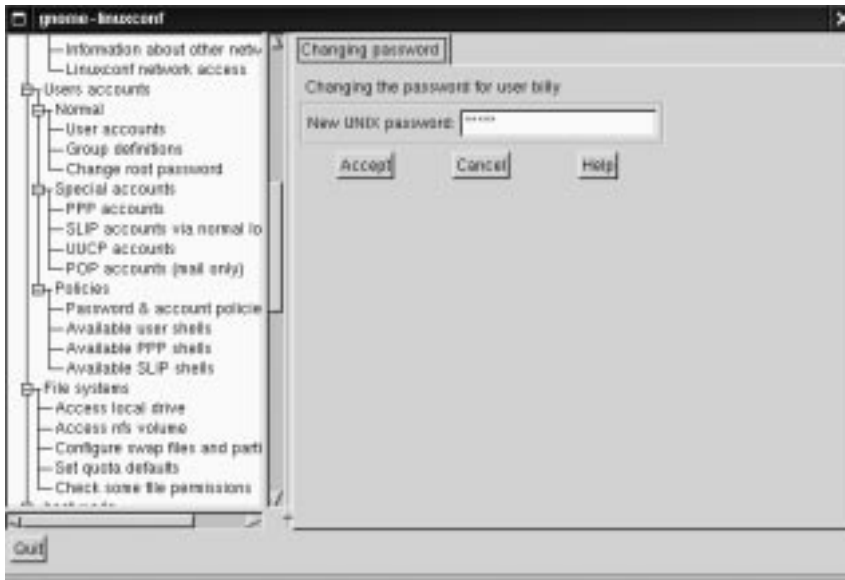
For both your root and user account, your passwords should be unique and easy enough for you to remember. (Passwords can't protect very well if they're jotted down on a piece of paper and taped to the monitor!)

What's both unique and easy to remember? Sometimes, passwords which contain both numbers and letters. Here's an example:

- Weak passwords: **airplane**, **icecream**, **california**

- Better passwords: **a!rpl8ne**, **!cec73am**, **c8Li7Qrnia**

One more thing: Passwords should be at least six characters in length -- upper and lowercase letters and/or numbers. (If you enabled **MD5 Passwords** during installation, your password can be up to 256 characters.)

Once you decide on a password you think you'll remember, type it in the box provided. You won't see the password, except in a series of asterisks (as shown in figure Figure 2–10, *Creating a User Account Password in Linuxconf*).

Then, click on the **Accept** button.

You'll be asked to retype the password for verification. Again, you won't see your password as you type it.

When you've finished, you'll see the account listed in the accounts panel (see Figure 2–11, *New User Account in Linuxconf* for an example).

That's all there is to it.

Figure 2–11   New User Account in Linuxconf
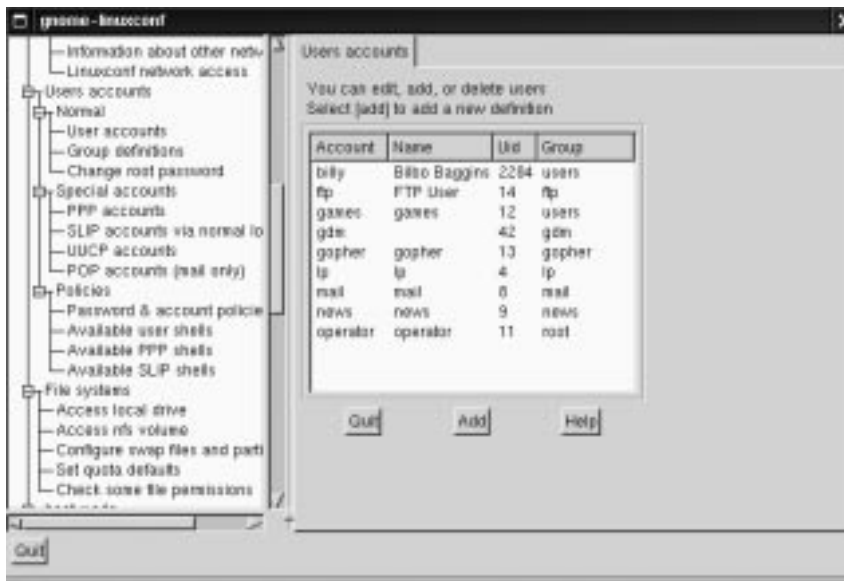
For additional security, you should change your passwords every now and then. From
your user account, you can change your password by clicking on the account name,
then clicking on the box marked **Passwd** at the bottom.

Tip

Now that you've created your user account, you might want
to reconsider whether your root account's password is se-
cure enough. You can change this password easily, from
within Linuxconf, by clicking on the **Change root password** item.

To learn more about how to modify your account or perform other account proce-
dures, turn to the *System Configuration* chapter in the *Official Red Hat Linux Refer-
ence Guide*.

From the shell prompt:

At the prompt, type useradd billy:

```
[root@localhost root]# useradd billy
[root@localhost root]#
```

Figure 2–12   Adding a User at the Shell Prompt
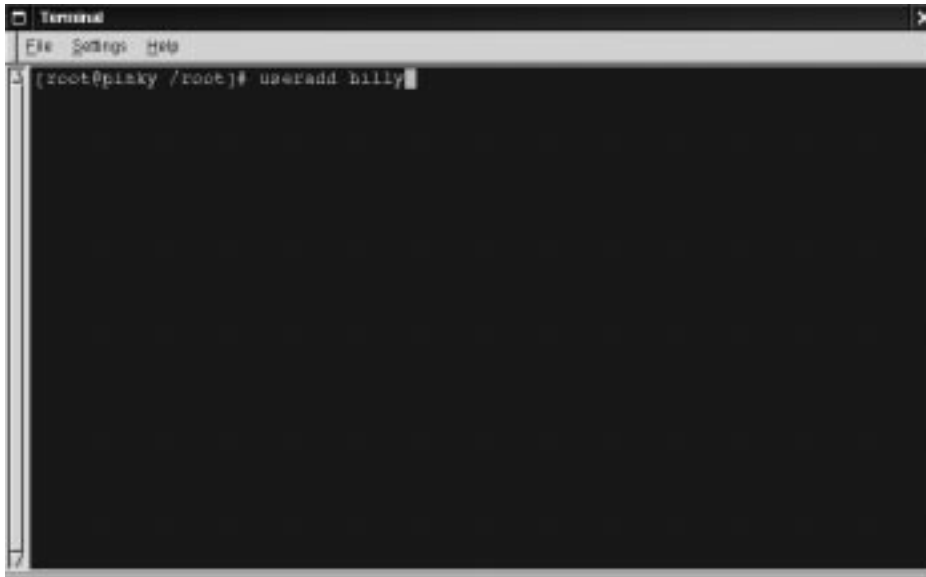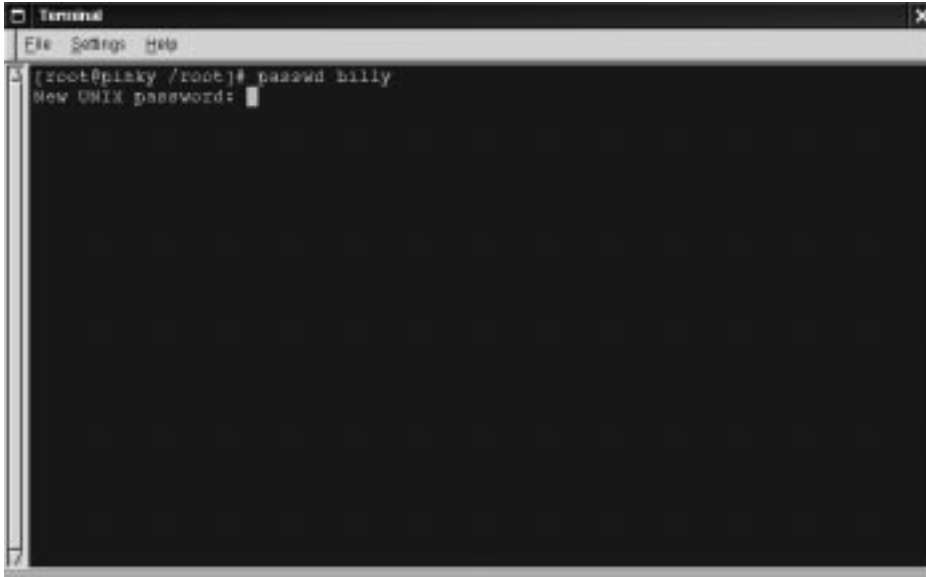
Looks like nothing's changed, right? Wrong. Although you've got the same prompt, an entry has already been made for the new account. Now, it's time to specify a password.

At the prompt, type `passwd billy`. It should look like this:

```
[root@localhost root]# passwd billy
[root@localhost root]#
```

Remember that the password should be easy to recall and a unique mixture of letters, symbols and/or numbers.

That's it.

Figure 2–13   Adding a Password at the Shell Prompt

From now on, whenever you want to add a user, change account information or change account passwords -- including the root password -- either `useradd` or Linuxconf will fit the bill. Make sure you're logged in as the root user, however.

Tip

If you find yourself switching around frequently between the root account and your user account, it can become confusing to know which account you're actually logged into. You can always tell you're in the root account you see the word [root at the start of the shell prompt or the hash mark (#) at the end. If you see (for example) [billy, or the dollar sign ($), you're working in your user account.

Regardless of the method you choose -- from the shell prompt or from Linuxconf -- your new user account's "login directory" will be placed in a subdirectory of /home.

To finish up and try out your new account, log out from your root account.

You'll be taken back to the login screen.

Now, you can log in to your new user account.

Summary

From X -- In a terminal emulation window, type linux-conf. Scroll to **User Accounts** => **Normal** => **User Accounts**. Click on **Add**. From the shell prompt -- type useradd *accountname*; at the shell prompt again, type passwd *accountname*. Type and re-type password.

2.7 Shutting Down

Some day, computers will probably be as easy to use as televisions are today.

At present, though, you can't simply turn off your computer when you're finished.
You can always log out from your account, which will return you to your login screen,
but if you want to completely shut off your machine, you've got a couple more steps
to take.

Figure 2–14   The Command `top` Shows Running Processes

Here's why: Even though you may not be typing, listening to music or browsing with
Netscape, your machine is still working on a variety of processes in the background.
(A **process** is a program which is being executed. Multiple processes are running all
the time on your system.)

Tip

Curious to take a peek? Just go to a shell prompt and type `top`. You'll see the processes that are currently running (see Figure 2–14, *The Command `top` Shows Running Processes*). To quit this view, type [Q].

(To learn more about the top command, type `man top` at the shell prompt; to move forward a screen, press the [Spacebar]; to move back a screen, press [B]; to quit, press [Q]. You'll learn more about these **man pages** later in this chapter.)

Like a faithful assistant, your Linux system is carrying out tasks silently all the time. You can't just turn out the lights and lock the door on your assistant. Instead, you've got to give them time to put away their work and make sure everything's in its proper place before saying "good night."

You can shutdown or reboot in a couple of ways while you're in GNOME:

From your GNOME session: Click on the **Main Menu Button**, choose **Log out**, select **Logout**, **Halt** or **Reboot**. You can also check **Save current setup**, which will preserve your desktop and running applications for the next time you log in. Before proceeding, you'll be asked to verify your decision by typing in your password. (If you're in your user account, you can type in your user password.)

From the graphical login screen: Go to **Options…** => **System** => **Halt…** or **Reboot…** You'll be asked whether you're sure you want to halt or reboot the machine. Click on **Yes** to proceed (a password isn't required).

Figure 2–15   The `shutdown` Command at the Shell Prompt
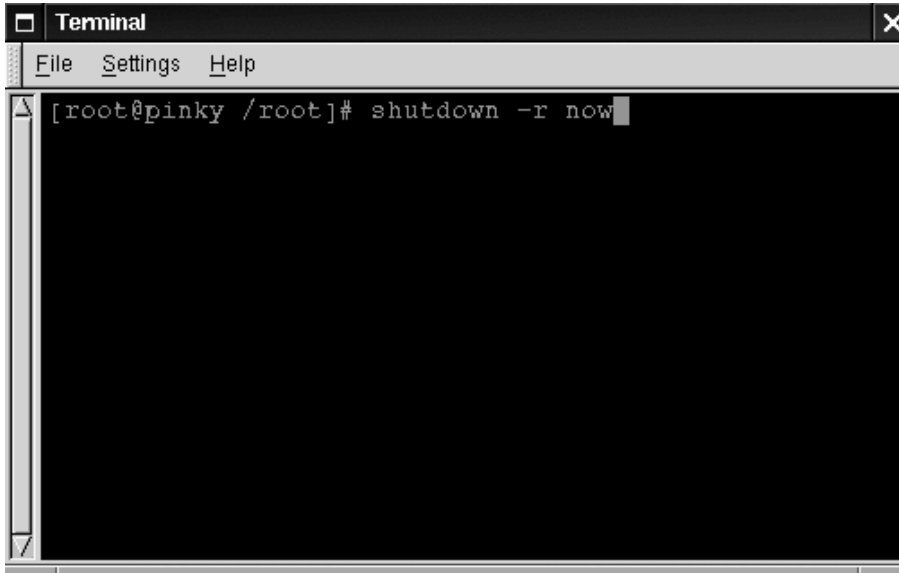
You can also reboot or halt your system from the shell prompt. In both cases, you'll be asked to provide your password to verify that you want to proceed with your decision.

To reboot from the prompt, type:

```
shutdown -r now
```

(See Figure 2–15, *The `shutdown` Command at the Shell Prompt.*)

Or, if you want to exit from your system and turn off your machine, type:

```
shutdown -h now
```

The `-r` option stands for "reboot," while the `-h` option means "halt." Stating `now` means that you want to perform this action immediately.

Please Note

Remember to save your work and exit from any applications which may be running before you perform a shutdown from the shell prompt, because you could lose work.

If you choose to halt the system, you'll see a list of messages about which services are stopping; then, you'll see a message similar to:

```
Power down
```

If your computer's BIOS supports the feature, you won't even have to turn off your machine; Red Hat Linux will simply "power down" the system. Now everything's put away and it's safe to turn off your computer.

Tip

Try substituting `+5` for `now`; you'll find that you've just commanded your assistant to put everything away and stop working in five minutes.

You can learn more about the shutdown command by typing:

```
man shutdown
```

at a shell prompt. You'll be presented with a "man page," which will tell you about this command.

To go forward a screen, press the [Spacebar]; to go back a screen, press [B]; and to quit, press [Q].

Summary

To shutdown or reboot from GNOME, from the login screen, go to **Options** => **System** => **Halt** or **Reboot**. From within your X session, go to **Main Menu Button** => **Log out**, **Reboot** or **Halt**. From the shell prompt: Log in as root, and type `shutdown -r now` (to reboot) or `shutdown -h now` (to halt).

2.8 Pulling Yourself Up by the Boot

When you're logged in as root, you might want to take a few minutes to create a fresh **boot disk** or copy the boot disk you already have.

There are a number of reasons you should make a boot disk: it can help you recover from a system failure, it can help you test a new kernel you've downloaded and compiled and it can help you share your computer with more than one operating system.

Tip

You can always use a copy of the boot disk to form the first half of a rescue disk set. You'll need a boot disk and a rescue disk to enter rescue mode. You can learn more about rescue mode techniques in the *Official Red Hat Linux Reference Guide*.

You were given the opportunity to make a boot disk when you installed Red Hat Linux. If you chose not to make a boot disk at installation, here's your chance to start from scratch.

For now, we'll make boot disks from the shell prompt while we're in an X session.

Go to the shell prompt: In GNOME, for example, click on the **Main Menu Button** on the panel, go to **Utilities** in the menu and click on one of the items marked **Xterm** or on the **GNOME terminal** item.

Now, make sure you're logged in as root. At the shell prompt, if you see something like [billy@localhost billy]$, for example, type:

```
[billy@localhost billy]$su -
Password: yourrootpassword
[root@localhost billy]#
```

This will allow you to change from your user account to the **root login** account.

Please Note

When you're root, you are "commander in chief" of everything on your system, so be careful when you issue commands.

Tip

The command su means substitute users, which lets you temporarily log in as another user. When you type su all by itself and press [Enter], you become root, or "superuser" while still inside your login shell. Typing su – makes you become root with root's login shell -- that is, it's as if you had logged in as root. What's the difference? There are some commands which can be performed only if you've logged in as root; so typing su – allows you to perform these commands without logging out and logging back in.

Briefly, we'll find the Linux kernel version; then, we'll use the mkbootdisk command to make a bootable floppy from the kernel.

Put a standard diskette in the floppy drive.

Tip

In Linux, the floppy drive is referred to as /dev/fd0.

If you've previously used the diskette, remember: You will lose anything that had been on that diskette!

Figure 2–16　Changing Directories to `/lib/modules`
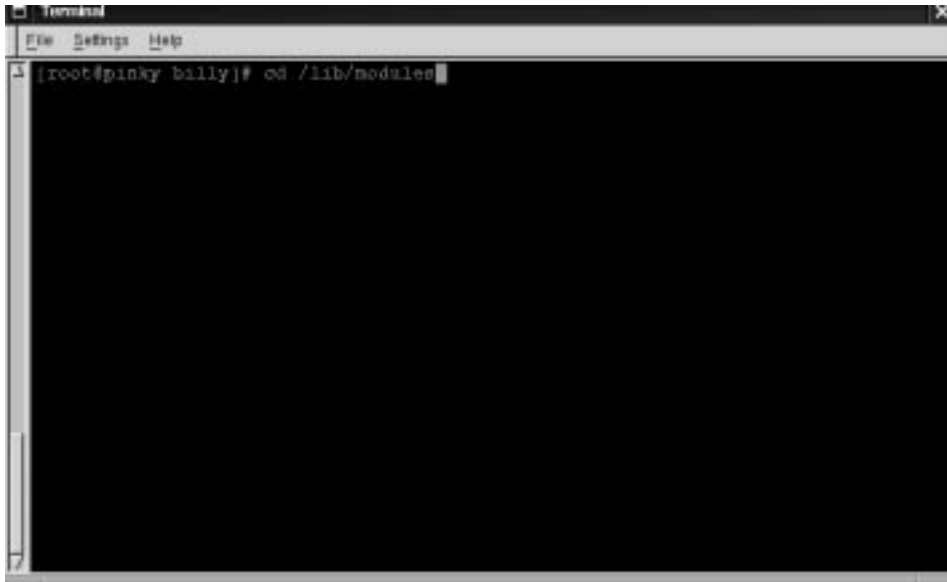
Now, at the prompt, type:

```
cd /lib/modules
```

(Shown in Figure 2–16, *Changing Directories to `/lib/modules`.*)

Now, type:

```
ls
```

The command `ls` will list the contents of a directory. (You can learn more about the `cd` and `ls` commands in Chapter 3, *You Are Here*. For now, just type the commands as you see them.)

Here, you can find the kernel version of your Red Hat Linux system. The kernel is the heart of any Linux system. Your kernel version will be something similar to:

```
2.2.x-yy
```

(there will be several numbers after `2.2`, as in `2.2.11-4`).

Now that you've found the kernel version, you can tell the `mkbootdisk` command which kernel to copy to your floppy. (If you don't tell `mkbootdisk` where to copy the kernel, it will default to copying to the floppy in `/dev/fd0`.)

Just type:

```
mkbootdisk --device /dev/fd0 2.2.x-yy
```

Then press [Enter].

Tip

If your screen becomes a little crowded with commands and "command not founds", you can always start with a clean slate by typing `clear` at the prompt.

You're done.

Summary

As root, in a terminal window, `cd /lib/modules`; choose kernel number; then type `mkbootdisk --device /dev/fd0 kernel.number`.

To clean the display, type `clear`.

2.9 A Good "Man" Is Easy to Find

As you investigate your new system, you're bound to have questions about commands and system services. One of the easiest ways to find out about how to use many commands or some applications is through the `man` command.

The word `man` stands for "manual," a series of online "pages" which can tell you the purpose of many commands. In a highly condensed format, man pages provide a summary of a command's purpose, the options available and the syntax which is used to issue the command.

If you count yourself among the "newbies" to Linux, you might not find man pages as useful as someone who's more accustomed to their terse delivery of information. But man pages can help steer you toward the proper way to use commands on your system. Even at this point, you can gain insight into your system by familiarizing yourself with the man pages. You'll certainly want to know how to use them eventually.

There are several ways to view the man pages: from GNOME's Help Browser, from an application called `xman` or from the shell prompt.

Figure 2–17   Man Page Index in GNOME's Help Browser

GNOME's Help browser:

To start GNOME's Help Browser, go to **Panel** => **Help system**. The GNOME Help Browser will start. On the first page, you'll find links to the man pages and to other helpful documentation.

Tip

Read about how to get more from GNOME in the GNOME User's Guide in this manual.

From `xman`:

Depending on your window manager, there will be different ways to access the graphical presentation of the man pages through the menus. (Window managers literally manage how the windows in your X session are presented.)

A quick way to start the one man page browser, however, is to go to a shell prompt and type:

```
xman
```

which will start one of the X Window System man page browsers. When the menu appears, click on **Manual Page**. From here, you'll have a number of options from which to find a man page: You can alphabetically display all the man pages on your system; search for them by command or by the section in which they appear.
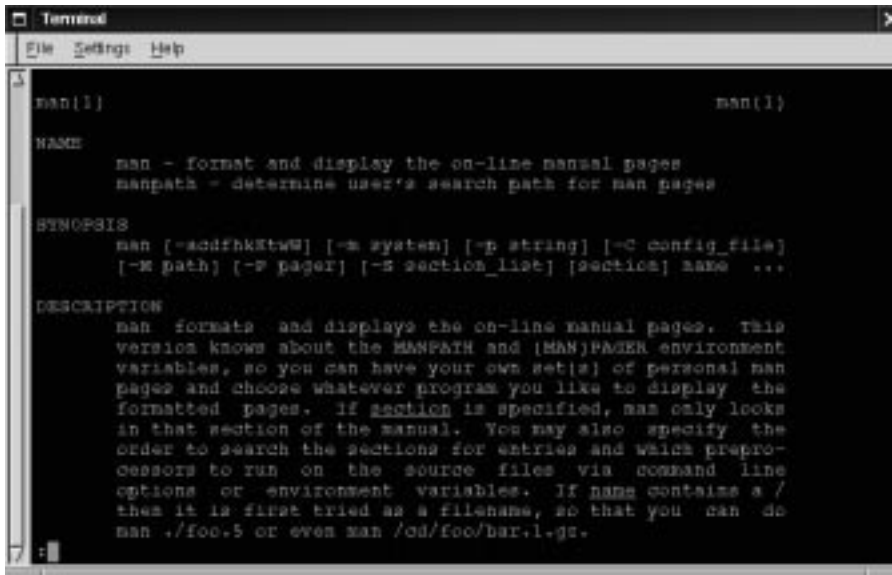
To see the scope of help available, bring up your manual browser, then click on **Manual Page** then, under **Options**, click on **Display Directory**. Here, you'll find the complete list of man pages available to you.

From the shell prompt:

If you're not in an X session, you can still read the man pages by typing

```
man pagename
```

at the shell prompt. To scroll forward through the document, press [Space]; to scroll back, press [B]. To quit the document, press [Q].

Figure 2–18   Reading a Man Page at the Shell Prompt

Of course, like any good help system, the man command has its own man page. At
the prompt, type

```
man man
```

to display the manual page (as shown in Figure 2–18, *Reading a Man Page at the
Shell Prompt*).

If you want to print:

Sometimes, just viewing the man page on the screen isn't enough; you may want to
have a printed copy in front of you. Although you could send the page to a printer,

because of certain text formatting in the man pages, you could end up with a document filled with "garbage," symbols which didn't translate from your screen to the printer.

Before you print, then, you may have to "strip" the formatting from the page, which you can do with the `col` command. (As you might guess, there's a man page for `col`, also.)

For example, to print a man page for the `man`, type:

```
man man | col -b | lpr
```

In detail, the above command "sends" the output of the manual page entry through the `col` filter, which helps format the output for the printer. Then, the output from `col` is sent to the printer. This is called **piping**, and you can learn more about it in Chapter 3, *You Are Here*.

3   You Are Here

Let's say you want to buy a pair of sneakers at a nearby shopping mall. You may not be familiar with the mall, but that shouldn't be a problem.

Why not? Because of the maps, which you can usually find near all the major entrances to the mall.

The same can be said for your Red Hat Linux system: Navigation's easy once you know where you're going.

Tip

Make sure you've logged into your user account! Remember that unless you like to live on the wild - and dangerous - side, using a root account for all your activities is toying with disaster. If you didn't create your user account at the time of your Red Hat Linux installation and haven't done so now, please do so now. Turn to Chapter 2, *Welcome to Linux* if you require assistance on how to create your user account.

3.1 Finding Yourself with `pwd`

Sooner or later (probably sooner), when you start looking through directories, you're bound to ask, "Where the heck am I?" And you won't be speaking philosophically.

DOS can answer that question just by showing you at the prompt like:

```
C:\GAMES\Game\ID1>
```

But Bash, the shell for your Linux system shows, by default, just shows your current directory.

Figure 3–1   The Command `pwd` Shows You Where You Are

```
[billy@pinky billy]$ pwd
/home/billy
[billy@pinky billy]$ ▮
```

Try this: open an Xterm window. You'll see something like:

```
[billy@localhost billy]$
```

Now type:

```
pwd
```

What do you see? Something like

```
/home/billy
```

The command pwd stands for **print working directory**. When you typed pwd, you were asking your Linux system, "where am I?" Your system responded by "printing" the directory you're in to the monitor -- also known as the **standard output**.

Seems easy, right? It ought to be; you'll be using pwd plenty as you look around. (Even Linux gurus can depend on this little command.)

3.2 Getting from Here to There: `cd`

Whenever you want to change directories, all you've got to do is type:

```
cd
```

Go ahead, try it in an Xterm window.

That didn't do much, did it? That's because you didn't tell your system where you wanted to go.

Whether you're going to a store in a mall or to visit relatives across the country, you've got to know how to get from one point to another. That is, you'll need to know the path to follow.

As with anything in life, the path -- or, pathname -- is basically the set of directions that takes you from one point to another. In the case of your Linux system (and in the DOS/Windows world, as well), you state a path to take you from one directory or file to another.

Let's try it again. Open an Xterm window. Find yourself first with the pwd command. When you type your commands, your window will look like:

```
[billy@localhost billy]$ pwd
/home/billy
[billy@localhost billy]$
```

Now that you see where you are, you can start to give your system the path to follow.

Well, almost…

Try typing:

```
cd home
```

What happened? You know there's a directory called home, and you typed in the path. So why does it say no such file or directory?

It means your path is incomplete.

Try typing:

```
cd /home
```

Now you've successfully changed directories and moved from your login directory into the subdirectory called `home`.

Figure 3–2   Absolute Pathnames State Full Path
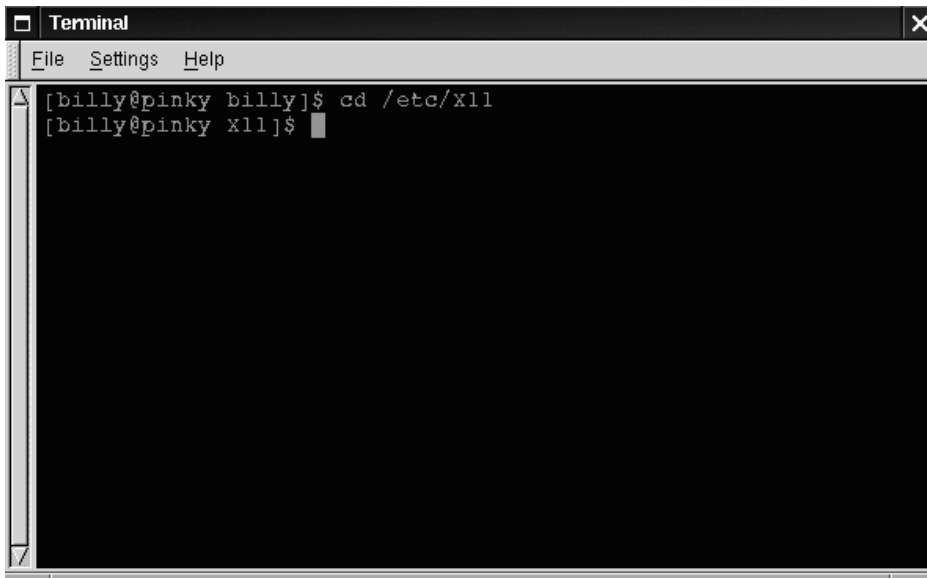
The difference, of course, was adding the forward slash.

Let's take a second to look at the reason adding a slash made all the difference.

When you saw you were in `/home/billy`, you were looking at the full path -- or the absolute path from the root directory. You can think of the `billy` directory as being located two directories "down" from the root, which is the topmost level of your system.

So when you typed:

```
cd /home
```

you were actually saying "go to the root directory, then go to the directory called home, which is one directory below the root." You specified an absolute path to get to the `home` directory.

Now, if you type:

```
cd /
```

you'll end up with a prompt that looks like:

```
[billy@localhost /]$
```

That single forward slash means you're at the root. When you're at the root, you can't go any higher on your system (the same is true in DOS/Windows).

To get back to your login directory from the root directory by using the absolute path, just type:

```
cd /home/billy
```

You're home.

Using the absolute path is just one way to move around. Another method of getting from one point to another is by using the relative path (as in Figure 3–3, *Relative Pathnames are Relative to Current Position*).

Let's go back to the `root` directory:

```
cd /
```

Now, let's move *back* to your login directory using relative pathnames:

```
cd home/billy
```
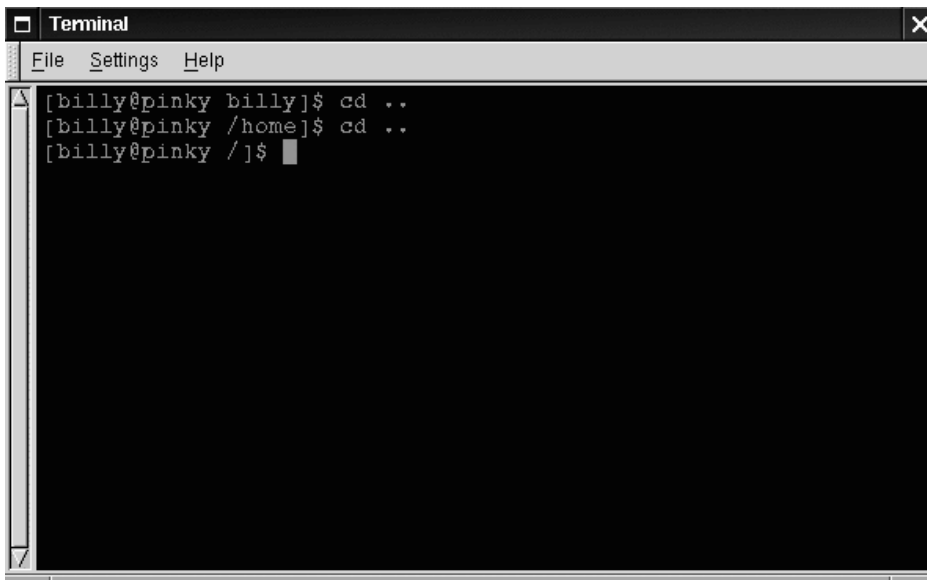
Notice that the first / is missing? That's because the root directory is the **parent** of the `home` directory, which means that the `home` directory is one step down from the root directory. Since `home` is the parent of the directory called `billy`, these two directories are separated with a /.

If you're in your login directory, you can move up one directory, to home, just by typing:

```
cd ..
```

The relative path describes the directory you want to `cd` to in terms which are relative to your current directory.

Figure 3–3   Relative Pathnames are Relative to Current Position

When you typed `cd  ..`, you were saying "go up one directory." The next directory up, from your login directory, was `home`.

Tip

When speaking of directories which hold other directories, you can refer to them as **parent directories**. In our case, `home` is the parent directory of `billy`.

Using two dots (`..`) when you `cd` is the same as stating you want to go to the parent
of your current working directory. Try using a single dot. Type:

```
cd .
```

What happened? Not much. That's because using a single dot (`.`) is the same as
specifying your current working directory.

The differences between absolute and relative paths can sometimes be pretty striking.

Getting back to our shopping mall analogy, if you were to give directions by using
an absolute path, you might say something like:

"Get the car keys. Get in the car. Start the car. Pull out of the driveway. Drive to the
corner…"

…And so on, until you're finally standing inside your favorite shoe store in the shop-
ping mall.

When you're using a relative path, you're saying something like:

"The store's a couple miles from here, in the shopping mall."

That's quite an exaggeration, but you get the idea: As long as you know where you
want to go in relation to where you are, you can use relative paths.

Tip

A path is absolute if the first character is a `/`; otherwise, it's
a relative path.

You're now in the `home` directory, the parent of your login directory. Type:

```
cd ..
```

and you'll find yourself at the root directory.

Using relative paths, get yourself back to your login directory by typing:

```
cd home/billy
```

Doesn't look much different from absolute paths, does it? Notice, though, that there's no forward slash in front of home. In essence, you were saying, "go one directory down, to home, then go to billy, in the home directory."

Tip

Whenever you want to quickly jump back to your login directory just type cd and press [Enter] anywhere in the system. You'll end up in your login directory.

That wasn't much of a demonstration.

Instead, from your login directory, type:

```
cd ../../etc/X11
```

Now, you're in the directory X11, which is where you'll find configuration files and directories related to the X Window System.

Please Note

You can always type pwd to find out where you are in the directory tree. And you can get back to your login directory with the cd command.

Take a look at your last cd command. What you were really telling your system was, "go up to the parent directory, then up to that directory's parent directory (which is the root directory), then go to the etc directory and from there, to the X11 directory."

Using an absolute path would also get you to the X11 quickly. Type:

```
cd /etc/X11
```

and you're there.

Tip

Always make sure you know which working directory
you're in before you state the relative path to the directory
or file you want to get to. You don't have to worry about
your position in the filesystem, though, when you state the
absolute path to another directory or file.

Now that you're starting to get the hang of changing directories, see what happens
when you change to root's login directory.

```
cd /root
```
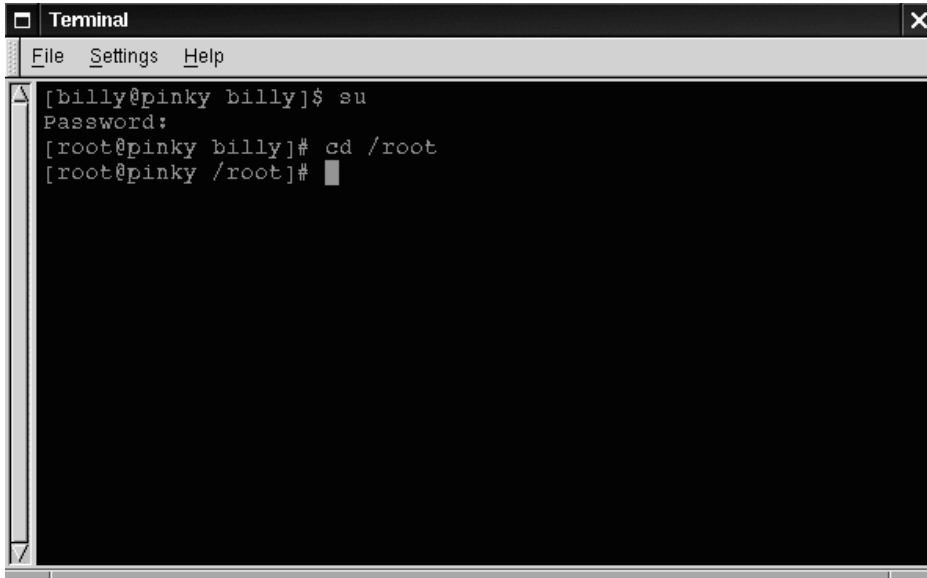
Oops… You're not logged in as root, so you're "denied permission" to access that
directory.

Denying access to the root and other users' accounts (or **login directories**) is one way
your Linux system prevents accidental or malicious tampering. You'll find out more
about file "ownership" and permissions later in this chapter.

Really want to change to the root login? Then you've got to use the su command.
Type this series of commands:

```
[billy@localhost billy]$ su
Password:your root password
[root@localhost billy]#cd /root
[root@localhost /root]#
```

As soon as you give the root password, you'll see the changes in your command
prompt to show your new, superuser status: the root account designation at the front
of the prompt and "#" at the end (as shown in Figure 3–4, *Becoming Root*).

Figure 3–4   Becoming Root

```
[billy@pinky billy]$ su
Password:
[root@pinky billy]# cd /root
[root@pinky /root]# ▊
```

Now, when you `cd` to root's login directory, you'll be granted access.

When you're done working as root, just type `exit` at the prompt.

```
[root@localhost /root]# exit
exit
[billy@localhost billy]$
```

Summary

To change directories using absolute pathnames, type `cd`
`/directory/directory`; to change directories using
relative pathnames, type `cd` *`directory`* to move one
directory below, `cd` *`directory/directory`* to move
two directories below, etc.; to jump from anywhere on the
filesystem to your login directory, type `cd`; to change to
the parent of the directory you're in, type `cd` `..` Use `.` to
refer to your current directory.

3.3 Looking Around with `ls`

Now that you know how to move around, it's time to take a look at what's in the
directories.

But first, let's make sure you've got something you can look for in your login direc-
tory before we go any further.

You can start by creating an empty file. To do so, you can use a utility called `touch`
at the shell prompt. Try it; type:

```
touch foo.bar
```

Now, in your login directory, you've got an empty file called `foo.bar`. You'll see
it in a couple minutes.

Let's also create a new directory, using the `mkdir` command.

At the prompt, type:

```
mkdir tigger
```

Now, you've created a directory called `tigger` in your login directory. From root,
your new directory's absolute pathname would be `/home/yourlogin/tigger`,
and your directory is the parent of `tigger`. (You can learn more about creating --
and removing -- files and directories in Chapter 4, *Managing Files and Directories*.)

Now, you're ready to go.

In the DOS world, using the `dir` command will display the contents of a directory.

The same can be said of Linux -- with some notable exceptions.

In Linux, `dir` won't fully display the contents of directories, and doesn't have the power or flexibility of the list command -- `ls`.

In your login directory, for example, type:
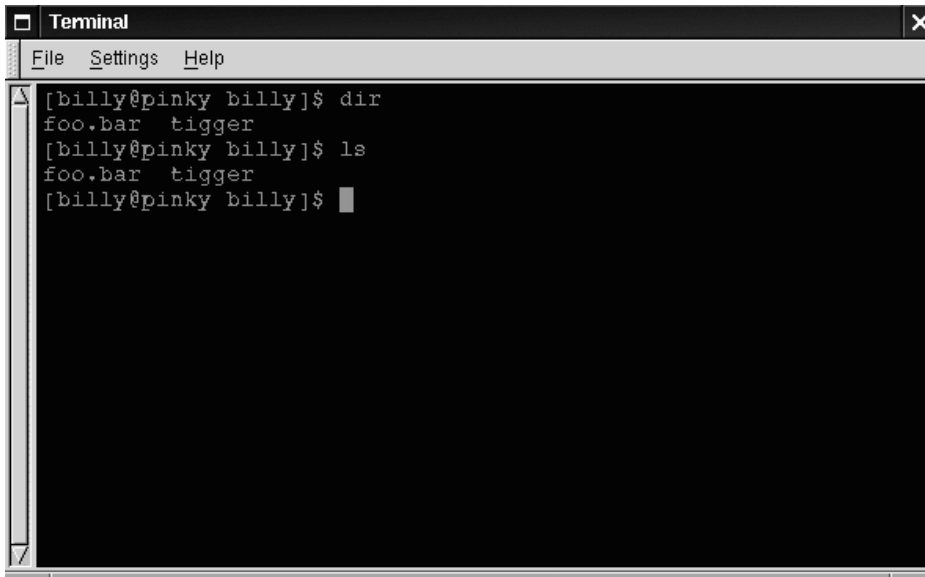
```
dir
```

Now, in the same Xterm window, type:

```
ls
```

Looks the same (see Figure 3–5, *dir and ls Seem Similar*). You see, among other contents, your new file, `foo.bar` and the new directory, `tigger`.

Figure 3–5 `dir` **and** `ls` **Seem Similar**

But here the similarities end. Where `dir` shows you the contents of your directory, it doesn't actually show you everything. Even using the `ls` command, by itself, won't show you all the files in your directory. To see everything, you've got to call upon an option or two.

For example, in the same window that you'd used to type the `dir` and `ls` commands, now type:

```
ls -a
```

Quite a difference. When you added the `-a` option, you were specifying that you wanted to list *all* the files in the directory (see Figure 3–6, `ls with -a` *Option*).

In fact, there are a multitude of options available with the `ls` command.

Tip

If you want to see all the options of the `ls` command, you can read the man page by typing `man ls` at a shell prompt. If you want to print the man page, type `man ls | col -b | lpr` at the prompt.

Why so many options? Because they can help you sort information according to your needs. For example, you can specify how files are displayed, see their **permissions** and much more.

———

Figure 3–6   `ls` with `-a` Option

```
┌─┬───────────────────────────────────────────────────────────┬───┐
│ □ │  Terminal                                                 │ × │
├───┴───────────────────────────────────────────────────────────┴───┤
│  File   Settings   Help                                            │
├────────────────────────────────────────────────────────────────────┤
│△│[billy@pinky billy]$ ls -a                                        │
│ │.                   .bash_history   .bashrc                       │
│ │..                  .bash_logout    foo.bar                       │
│ │.Xdefaults          .bash_profile   tigger                        │
│ │[billy@pinky billy]$ █                                            │
│ │                                                                  │
│ │                                                                  │
│ │                                                                  │
│ │                                                                  │
│ │                                                                  │
│ │                                                                  │
│ │                                                                  │
│▽│                                                                  │
└────────────────────────────────────────────────────────────────────┘
```

When you typed `ls -a`, you probably noticed the files that begin with dots. These are called **hidden files** or, appropriately enough, **dot files**.

Hidden files are mostly configuration files which set preferences in programs, window managers, shells and more. The reason they're "hidden" is to help prevent any accidental tampering by the user. Besides, when you're searching for something in a directory, you're not often looking for these configuration files, so keeping them hidden helps to avoid some screen clutter..

Whenever a filename starts with a dot (.), it's a hidden file, and `ls` won't list it.

Viewing all the files can give you plenty of detail, but there's more you can uncover, simply by adding more than one option.

If we want to see the size of a file or directory, when it was created and more, we can just add the "long" option (`-l`) to our `ls -a` command.

Try it. Type:

```
ls -al
```

There's quite a bit more detail now. You can see the file creation date, its size, own-
ership, permissions and more.

You don't have to be in the directory whose contents you want to view, either.

Let's see what's in the /etc directory by typing:

```
ls -al /etc
```

Here, you'll get plenty of information about the contents of the /etc directory.

If you want to add color to your listing, just include the --color option.
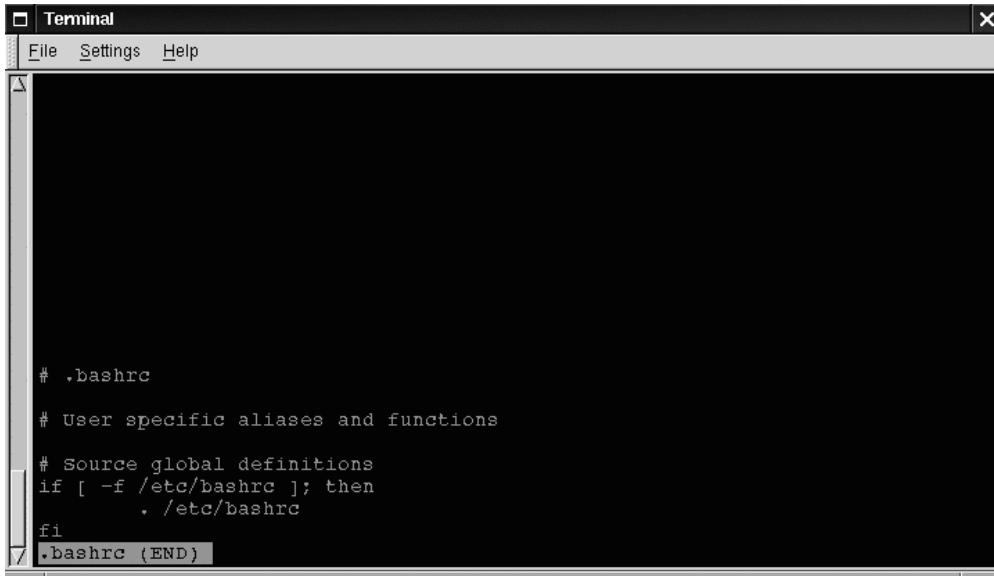
```
ls -al --color /etc
```

To some, adding --color does more than add a splash of color; it gives a clue about
the types of files in a directory. For example, directories might all be a royal blue,
program files would be green, and so on.

If you like what you see, here's how you can display the listing in color all the time.
Briefly, we'll be adding one line to the .bashrc file in our login directory.

The .bashrc file is used by your shell when you login (an example .bashrc file
is shown in Figure 3–7, *The .bashrc File*).

Now before we go any further...

Figure 3–7   The `.bashrc` File

```
┌─┬──────────────────────────────────────────────────────────────────┬─┐
│□│ Terminal                                                          │✕│
├─┴──────────────────────────────────────────────────────────────────┴─┤
│ │ File   Settings   Help                                             │
├─┬────────────────────────────────────────────────────────────────────┤
│∆│                                                                     │
│ │                                                                     │
│ │                                                                     │
│ │                                                                     │
│ │                                                                     │
│ │                                                                     │
│ │                                                                     │
│ │ # .bashrc                                                           │
│ │                                                                     │
│ │ # User specific aliases and functions                              │
│ │                                                                     │
│ │ # Source global definitions                                        │
│ │ if [ -f /etc/bashrc ]; then                                         │
│ │         . /etc/bashrc                                               │
│ │ fi                                                                  │
│∕│ .bashrc (END)                                                       │
└─┴─────────────────────────────────────────────────────────────────────┘
```

Remember that any changes you make to configuration files can cause you a world of grief if you've made a mistake and you don't have a backup copy of that file.

To make a backup copy, make sure you're in your login directory and in an Xterm window, type:

```
cd
```

to get to your login directory. Then copy the `.bashrc` file, keeping it in the same directory, but with a name like `.bashrc2`.

```
cp .bashrc .bashrc2
```

When you type the above command, what you're saying is, "make a copy of the `.bashrc` file and name that copy `.bashrc2`."

Now, you have a backup copy of the unmodified `.bashrc` file in your login directory. If you make a mistake or have trouble, you can replace your `.bashrc` file by typing:

```
cp .bashrc2 .bashrc
```

at the shell prompt.

If you need to type this command, you'll be saying, "make a copy of the file `.bashrc2` and name that copy `.bashrc`." The copy command here will overwrite the original `.bashrc` file -- and you'll still keep a copy of the original (and untouched) `.bashrc` file with the name of `.bashrc2`.

Now that we're prepared, we'll open `.bashrc` with Pico, a simplified **text editor**. (A text editor is a utility program that can create or modify files.) From an Xterm window, type:

```
pico .bashrc
```

You should see something like this:

```
# .bashrc

# User specific aliases and functions

# Source global definitions
if [ -f /etc/bashrc ]; then
        . /etc/bashrc

fi
```

It's a pretty short file. Those hash marks (`#`) are **comments**. Any text after them is ignored by the shell, but they are put there to help guide anyone who's editing or modifying files.

Bring your cursor under the line `#User specific aliases and functions` and type:

```
alias ls="ls -al --color"
```

So the full file ought to look something like this:

```
# .bashrc

# User specific aliases and functions

alias ls="ls -al --color"

# Source global definitions
if [ -f /etc/bashrc ]; then
        . /etc/bashrc

fi
```

See Figure 3–8, *Adding an Alias for* `ls` *to* `.bashrc` for an example in Pico.

Figure 3–8   Adding an Alias for `ls` to `.bashrc`

Double-check for any typos then, when you're satisfied with the changes, exit by pressing the [Ctrl] and [X] keys. You'll see, at the bottom of your editor screen, a message reading:

```
Save modified buffer (ANSWERING "No" WILL DESTROY CHANGES)?
```

Press [Y] for "yes." Now, another message will appear at the bottom:

```
File Name to write: .bashrc
```

Simply pressing [Enter] will save your changes to your `.bashrc` file.

You won't be able to see your changes take effect until you close your Xterm window and open a new Xterm. Once you do that, you'll see your modifications take effect.

Here's a short list of some popular options with `ls`. Remember, you can view the full list by reading the `ls` man page (`man ls`).

- `-a` -- all. Lists all the files in the directory, including the hidden files (`.file-name`). The `..` and `.` at the top of your list refer to the parent directory and the current directory, respectively.

- `-l` -- long. Lists details about contents, including permissions (modes), owner, group, size, creation date, whether the file is linked to somewhere else on the system and where its link points.

- `-F` -- file type. Adds a symbol to the end of each listing. These symbols include `/` to indicate a directory; `@` to indicate a symbolic link to another file; `*` to indicate an executable file.

- `-r` -- reverse. Lists the contents of the directory from back to front.

- `-R` -- recursive. This recursive option lists the contents of all directories (below the current directory) recursively.

- `-S` -- size. Sorts files by their size.

A little later in this chapter, when we introduce you to **pipes** and **I/O redirection**, you'll discover that there are other ways to view the contents of a directory.

Summary

To see the contents of a directory, type `ls` at a shell prompt;
typing `ls -a` will display all the contents of a directory;
typing `ls -a --color` will display all the contents cat-
egorized by color.

3.4 A Larger Picture of the Filesystem

Every operating system has a method of storing its files and directories so that it can keep track of additions, modifications and other changes.

In Linux, every file is stored on the system with a unique name, in directories which can also hold other files and directories -- or, subdirectories.

You might think of the system as a tree-like structure, in which directories "branch off." Those directories may contain -- or be the "parent" of -- other directories which may hold files or directories of their own.

There wouldn't be a tree without a root, and the same is true for the Linux filesystem. No matter how far away the branches, everything is connected to the root, which is represented as a single forward slash (/).

It might seem confusing to have several references to "root" - the root account, the root account's login directory and the root directory (/), but think of it this way: The root login, who is the system administrator, is just as important to keeping things together in the system as the system's root (/).

Tip

Even though there are other Linux distributions, your Red Hat Linux system is likely to be compatible with them. The reason is because of the **Filesystem Hierarchy Standard** (also known as FHS). These guidelines help to standardize the way system programs and files are stored on all Linux systems.

To read more about the FHS, turn to the chapter on system administration in the Red Hat Linux Reference Guide. You can also visit the FHS website: http://www.path-name.com/fhs.

As long as we're logged into our user account -- which will help prevent disastrous mistakes -- let's take a look around.

The first stop on this tour ought to be the root directory, which will give us a larger picture of where things are.

At the shell prompt, then, let's type:

```
cd /
```
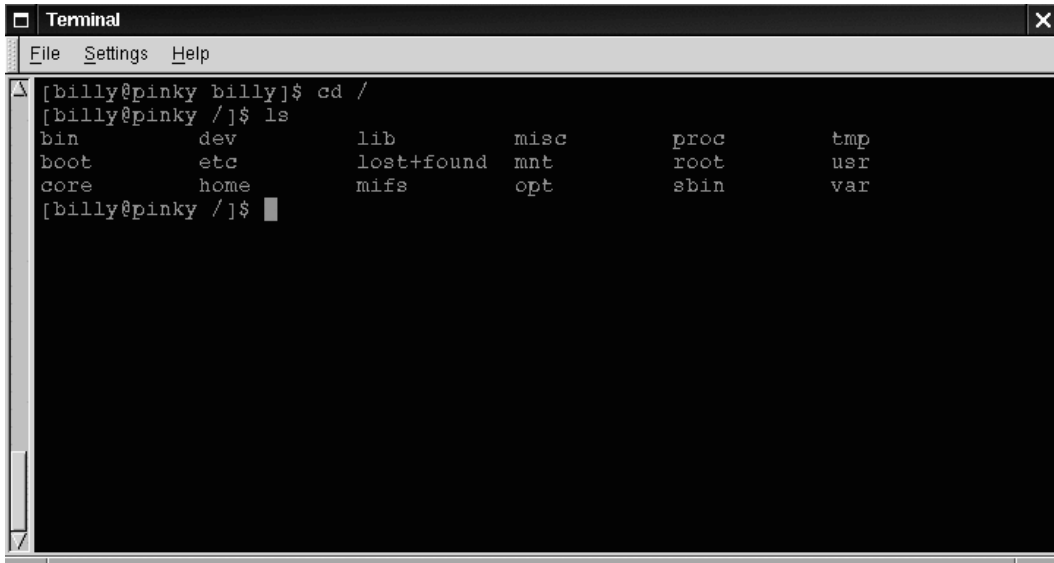
We'll see a prompt which looks like:

```
[billy@localhost /]$
```

Now, let's see which directories "branch off" root by typing:

```
ls
```

Doesn't look like much, does it? Well, it's a little like viewing the tip of an iceberg. These are the parent directories of other directories, in which there may be other directories… and so on.

Figure 3–9   A View of the Directories from Root

```
[billy@pinky billy]$ cd /
[billy@pinky /]$ ls
bin           dev           lib           misc          proc          tmp
boot          etc           lost+found    mnt           root          usr
core          home          mifs          opt           sbin          var
[billy@pinky /]$ ▊
```

Here are just a few of the directories we're likely to find:

```
etc      lib      sbin
usr      var
```

There are more, but let's take a look in the /etc directory.

```
[billy@localhost /]$ cd etc
[billy@localhost /etc]$ ls
```

Here, among other types of files and directories, we'll find **configuration files**, which are files that help make programs work for our system, store our program and system settings and more.

Among the directories in here, you'll see /etc/X11, which also contains directories and configuration files for the X Window System.

In the directory /etc/skel, you'll find **skeleton** user files, which are used to populate newly created user accounts with standard, commonly used files.

That sounds a little gothic, perhaps, but here's what it means. When we were logged in as root, one of the first tasks we undertook was to create a user account.
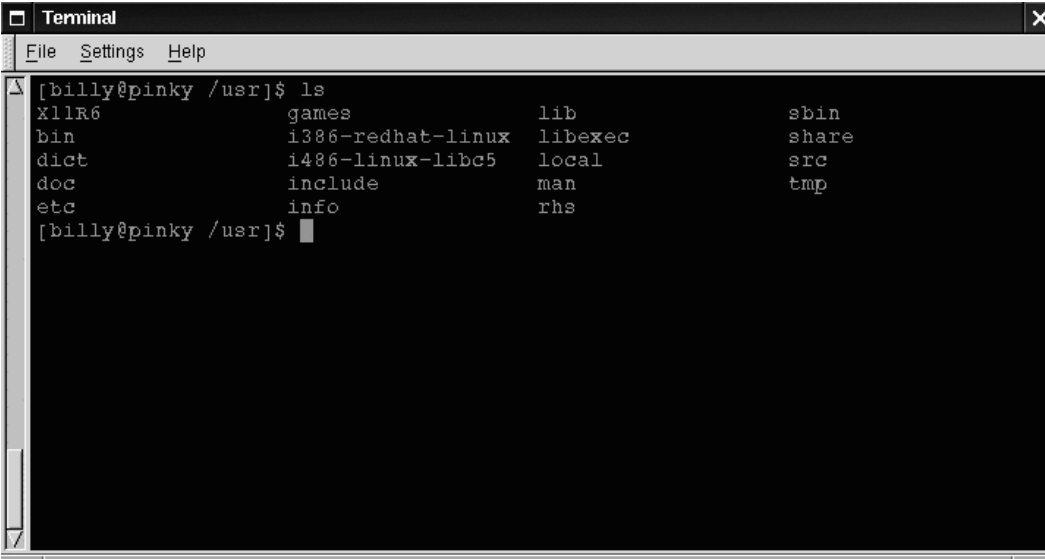
When that user account was created, files were taken from /etc/skel and placed into the new account. These files helped to "flesh out" the account (skeleton… flesh… get it?).

Let's look around a little in /usr. From our current location in /etc/skel, we can type:

```
[billy@localhost /skel]$ cd /usr
[billy@localhost /usr]$ ls
```

In /usr, we're going to find a number of directories which hold some of our system's most important programs and files (see Figure 3–10, *Typing ls in /usr*).

Figure 3–10   Typing `ls` in `/usr`

In /usr/man we'll find the system manual pages; other documentation which isn't covered by man pages will be found in /usr/doc and in /usr/info.

In /usr/X11R6, we'll find files related to the X Window System, including configuration and documentation files.

Although we may think of something more literary when we hear the word "libraries," in /usr/lib we'll find files which are considered **libraries** for our system. In this context, libraries are files containing commonly used snippets of code which can be shared by many programs.

Red Hat Linux uses the RPM (the *R*ed Hat *P*ackage *M*anager) technology of software installation and upgrades. Using RPM, either from the shell prompt or through Gnome-RPM, is both a safe and convenient way to upgrade or install software.

(For more on using Gnome-RPM, see Chapter 6, *Using Gnome-RPM and RPM*, or turn to its chapter in the Red Hat Linux Reference Guide.)

However, once you become more comfortable with your system, there may be times when you'll want to install software that may not be available in RPM format. To minimize collisions with RPM-managed files, the best place to put such software is in /usr/local.

3.5 "Washing" the Window

After even one ls command in an Xterm window, things might start feeling a little crowded. We can always exit from the terminal window and open a new one, but here's a quicker way to wipe the slate clean.

Just type:

```
clear
```

at the shell prompt. The clear command does just as advertised: it clears the terminal screen.

Sometimes, you may accidentally open a program file or some other non-text file in a terminal window. Once you close the file, you could find that the text you're typing doesn't match with the output on the monitor.

In such cases, you simply have to type:

```
reset
```

to return the window to its default values.

Summary

To clear clutter in a console or Xterm window type `clear`;
to return an Xterm window to its default display properties,
type `reset`.

3.6 Using `cat`

There's a handy utility which can help you keep short lists, gather those lists together
and, at the same time, show you a little of the power behind your Red Hat Linux
system.

The utility is called `cat`, short for "concatenate," which means that it strings files
together.

But `cat` can also perform a quick demonstration of two important terms: **standard
input** and **standard output**.

Standard input and standard output direct input and output (often referred to as **I/O**)
to the user. If a program reads from standard input, it will, by default be reading from
the keyboard. If a program writes to standard output, by default it will be writing to
the screen.

Start `cat` to see what this means. At the shell prompt, type:
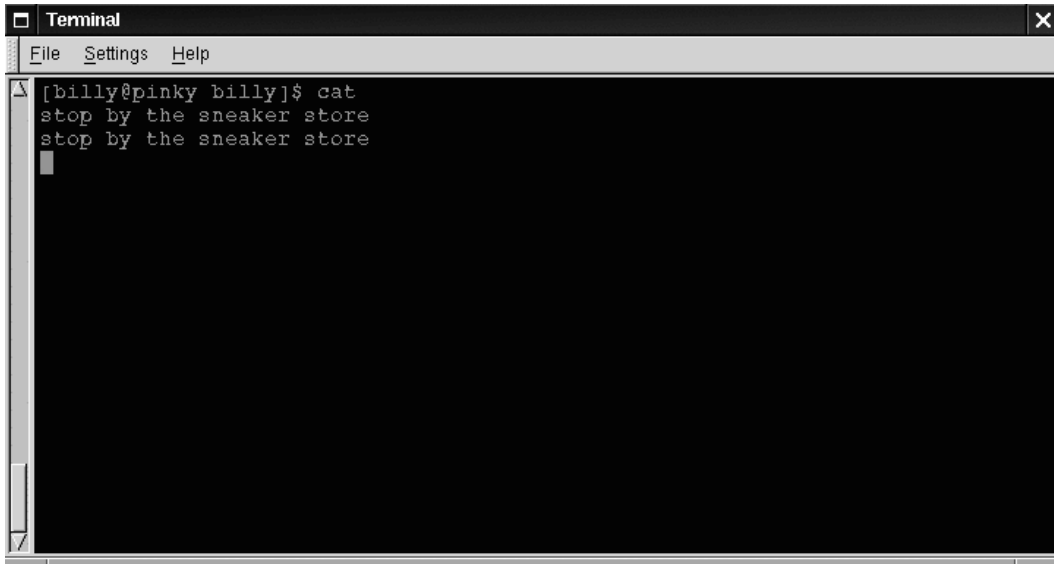
```
cat
```

The cursor moves to a blank line. Now, in that blank line, let's type:

```
stop by sneaker store
```

and press the [Enter] key. Your screen will look like:

```
[billy@localhost billy]$ cat
stop by sneaker store
stop by sneaker store
```

Figure 3–11 `cat` **Demonstrates Standard Input and Standard Output**

To quit `cat` now, just move the cursor to a blank line by pressing [Enter] then press the [Ctrl] and [D] keys at the same time.

So it's not too exciting. But `cat` has just demonstrated the definition of standard input and standard output. Your input was read from the keyboard (standard input), and that input was then directed to your terminal (standard output).

Summary

Standard input is often text which is entered from the keyboard. Standard output is the place where information is sent, such as your terminal (as shown in Figure 3–11, *cat Demonstrates Standard Input and Standard Output*).

3.7 Using Redirection

Now that we have a handle on what standard input and standard output are, it's time to expand a little.
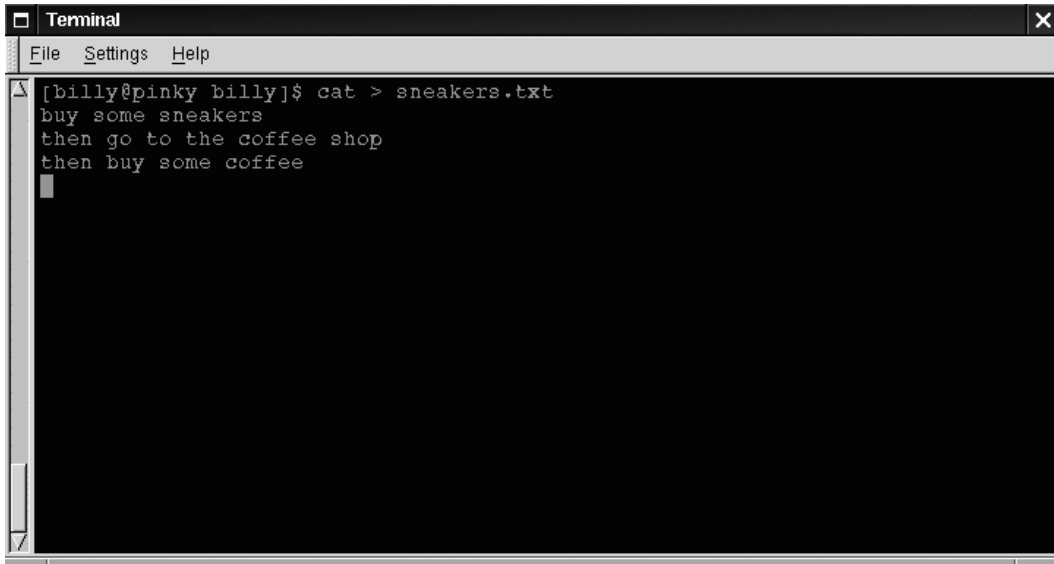
Redirection means causing the shell to change what it considers standard input or where the standard output is going.

We used cat before to demonstrate the idea behind standard input and standard output. Now, let's use cat to see how standard output can be redirected.

To redirect standard output, we'll use the > symbol. Placing > after the cat command (or after any utility or application that writes to standard output) will direct its output to the filename following the symbol.

Let's try it. In an Xterm window type:

```
[billy@localhost billy]$ cat > sneakers.txt
buy some sneakers
then go to the coffee shop
then buy some coffee
```

Figure 3–12   Redirecting Output to a File

Now press [Enter] to go to an empty line, and use the [Ctrl]-[D] keys to quit `cat`.

Notice the difference (see Figure 3–12, *Redirecting Output to a File*)? For one thing, there are no double entries. That's because the standard output from cat was redirected. That redirection was to a brand new file you made called `sneakers.txt`.

You can find the file in your login directory (may we suggest using `ls` if you want to see it listed?).

You can even use `cat` to read the file, by typing:

```
cat sneakers.txt
```

at the prompt.

Tip

Be careful when you redirect the output to a file, because you can easily overwrite an existing file! Make sure the name of the file you're creating doesn't match the name of a pre-existing file, unless you want to replace it.

Let's use output redirection for another file and call it home.txt.

```
[billy@localhost billy]$ cat > home.txt
bring the coffee home
take off shoes
put on sneakers
make some coffee
relax!
```

Now, on an empty line, use the [Ctrl]-[D] keys again to quit cat.

We can check the file again by typing:
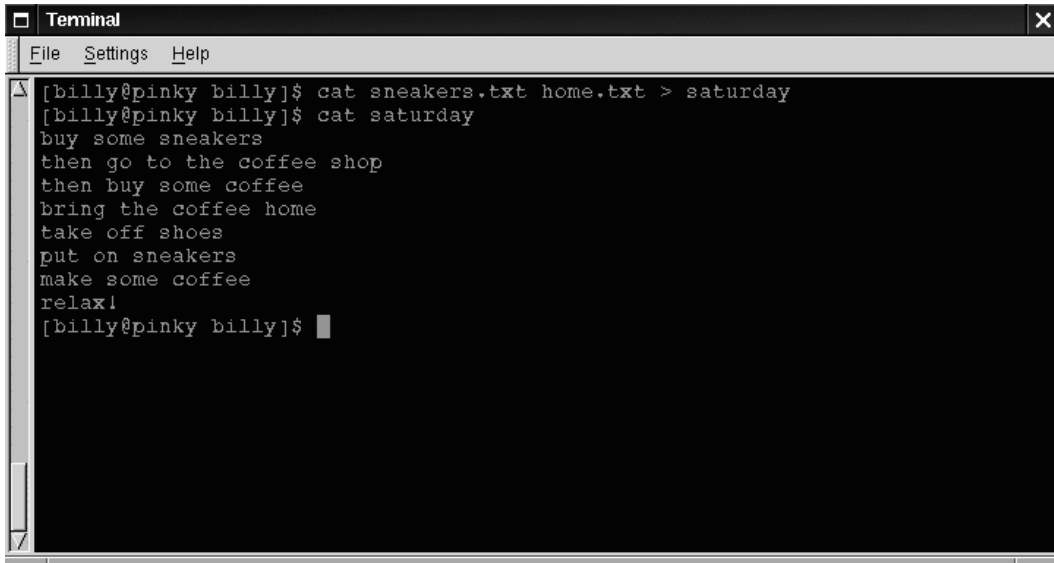
```
cat home.txt
```

at the prompt.

Let's use cat again to join home.txt with sneakers.txt and redirect the output of both files to a brand new file we'll call saturday (you'll find an example in Figure 3–13, *Joining Files and Redirecting Output*).

```
[billy@localhost billy]$ cat sneakers.txt home.txt > saturday
```

That's it.

Figure 3–13   Joining Files and Redirecting Output

```
┌──┬─────────────────────────────────────────────────────────────────────┬───┐
│ □ │ Terminal                                                          │ × │
├──┴─────────────────────────────────────────────────────────────────────┴───┤
│  File  Settings  Help                                                       │
├─────────────────────────────────────────────────────────────────────────┤
│ [billy@pinky billy]$ cat sneakers.txt home.txt > saturday                 │
│ [billy@pinky billy]$ cat saturday                                         │
│ buy some sneakers                                                         │
│ then go to the coffee shop                                                │
│ then buy some coffee                                                      │
│ bring the coffee home                                                     │
│ take off shoes                                                            │
│ put on sneakers                                                           │
│ make some coffee                                                          │
│ relax!                                                                    │
│ [billy@pinky billy]$ █                                                    │
│                                                                           │
│                                                                           │
│                                                                           │
└─────────────────────────────────────────────────────────────────────────┘
```

Now it's time to check our handiwork. Type:

```
[billy@localhost billy]$ cat saturday
```

and you should see something like this:

```
[billy @localhost billy]$ cat saturday
buy some sneakers
then go to the coffee shop
then buy some coffee
bring the coffee home
take off shoes
put on sneakers
make some coffee
relax!
[billy @localhost billy]$
```

You can see that cat has added home.txt where sneakers.txt left off.

Tip

Creating and combining short files with `cat` can be a convenient alternative to using a text editor like Pico.

Summary

By using the output redirection symbol (>) you can send the output to a file instead of the terminal. The `cat` utility can be used along with output redirection to join files together into a single, unified file with one filename.

3.8 Appending Standard Output

There's a neat twist to output redirection which allows you to add new information to the end of an existing file. Similar to when you used the > symbol, you tell your shell to send the information somewhere other than standard output.

However, when you use >>, you're *adding* information, rather than replacing it.

The best explanation is a demonstration, so let's take two files which have already been created -- `sneakers.txt` and `home.txt` -- and join them by using the append output symbol. We want to add the information in `home.txt` to the information already in `sneakers.txt`, so we type:

```
cat home.txt >> sneakers.txt
```

Now let's check the file by typing:

```
cat sneakers.txt
```

And there it is -- with the contents of `home.txt` at the end.

What we were saying when we typed that command was, "append the output from the file `home.txt` to the file `sneakers.txt`."

By appending the output, we've saved ourselves a step or two (and a bit of disk clutter) by using existing files, rather than creating a new file.

Compare the results of the files sneakers.txt and saturday now, and you'll see that they're identical. To make your comparison, just type:
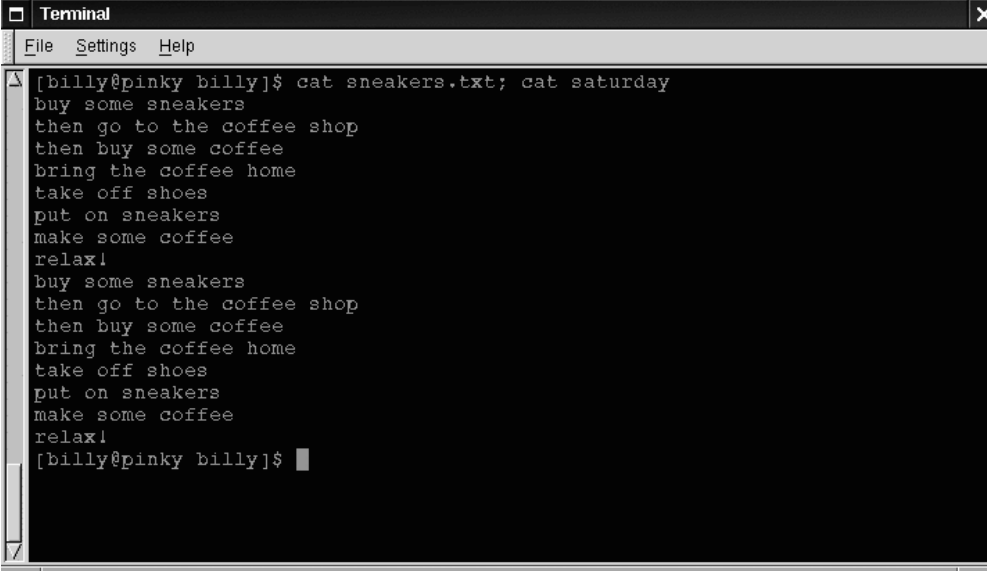
```
cat sneakers.txt; cat saturday
```

The contents of both files will be displayed - first sneakers.txt, then saturday (as shown in Figure 3–14, *Stringing Commands and Comparing Files*).

Tip

Remember that when you append output, you've got to include two greater-than symbols (>>). Otherwise, you'll end up replacing the very file to which you want to append information!

Figure 3–14   Stringing Commands and Comparing Files

```
[billy@pinky billy]$ cat sneakers.txt; cat saturday
buy some sneakers
then go to the coffee shop
then buy some coffee
bring the coffee home
take off shoes
put on sneakers
make some coffee
relax↓
buy some sneakers
then go to the coffee shop
then buy some coffee
bring the coffee home
take off shoes
put on sneakers
make some coffee
relax↓
[billy@pinky billy]$ ▊
```

(By the way, if you're curious about the use of the semi-colon in that last command, read on. We'll cover that later in this chapter.)

Summary

To append output, use two greater-than symbols (>>). For example: `cat addthisfile >> tothisfile`.

3.9 Redirecting Standard Input

Not only can you redirect standard output, you can perform the same type of redirection with standard input.

Here's how it works:

When you use the redirect standard input symbol <, you're telling the shell that you want a file to be read as input for a command.

We can use a file we've already created to demonstrate this idea. Just type:

```
cat < sneakers.txt
```

Because we used the less-than symbol (<) to separate the cat command from the file, the output of sneakers.txt was read by cat.

3.10 Pipes

No, we're not talking about plumbing here. In Linux, **pipes** connect the standard output of one command to the standard input of another command.

Let's take a step back, to the ls command. There are plenty of options available with ls, but what if the contents of a directory stream by too quickly for you to view them?

Let's view the contents of the /etc directory.

```
ls -al /etc
```

How do we take a closer look at the output before it races off the screen?

One answer is to pipe the output to a utility called less. Known as a **pager,** less, (like more) allows us to view information one page (or screen) at a time.

We use the vertical bar (|) to pipe the commands (as shown in Figure Figure 3–15, *Piping Output of ls to less*).

```
ls -al /etc | less
```

Now we can view the contents one screen at a time. To move forward a screen, just press [Space]; to move back a screen, press [B]; to quit, just press [Q].

Figure 3–15   Piping Output of `ls` to `less`

Actually, we've already been using pipes, before we even discussed what they were.

In previous references to man pages, we used the following to print man page entries:

```
man ls | col -b | lpr
```

Here, we're sending the output of man `ls` to a filter called `col` with an option of `-b` to help format the text for the printer, then we sent the output of that to the printer using the `lpr` command.

Summary

Piping allows you to send the output of one command as the input of another command. For example: `ls -al /etc | more` pipes the output of the `ls` command to the `more` utility for easy viewing.

3.11 Stringing Commands Together

Linux allows you to enter multiple commands at one time. The only requirement is that you separate the commands with a semicolon (you can see an example in Figure Figure 3–14, *Stringing Commands and Comparing Files*).

Want to see how long you've been online using Netscape? Just combine the date command with Netscape's command.

```
date; netscape; date
```

Remember that commands are case sensitive, so the command to start Netscape must be in lower-case to start the browser.

In the Xterm window, we'll see something like this:

```
[billy@localhost billy]$ date; netscape; date
Fri Jul 26 13:26:27 EST 1999
```

We'll see the second date entry when we close out of Netscape. Then, our screen will look like this:

```
[billy@localhost billy]$ date; netscape; date
Fri Jul 26 13:26:27 EST 1999
Fri Jul 26 14:28:32 EST 1999
[billy@localhost billy]$
```

And the prompt will return. The discrepancy between the two results from the date command shows that we were using Netscape for just over an hour.

3.12 Ownership and Permissions

Earlier in this chapter, when we tried to cd to root's login directory, we received the following friendly message:

```
[billy@localhost billy]$ cd /root
bash: /root: Permission denied
[billy@localhost billy]$
```

That was one demonstration of Linux's security features. Linux, like UNIX, is a multi-user system, and file permissions are one way the system uses to protect against any type of tampering -- malicious or accidental.

One way to gain entry when we see we're denied permission is to su to root, as we learned earlier. That's because whoever knows the root password has complete access.

```
[billy@localhost billy]$ su
Password: your root password
[root@localhost billy]# cd /root
[root@localhost /root]#
```

But switching to superuser isn't always convenient -- or smart, since it's so easy to mistakenly mess up important configuration files.

All files and directories are "owned" by the person who created them. We created the file sneakers.txt in our login directory, so sneakers.txt "belongs" to us.

That means, we can specify who's allowed to read the file, write to the file or, if it's an application instead of a text file, who can execute the file.

Reading, writing and executing are the three main settings in permissions.

Since every user on the system is placed into a **group** when that user is created, then we can also specify whether certain groups can read, write to, or execute our file.

Let's take a closer look at sneakers.txt with the ls command using the -l (long) option (see Figure 3–16, *Permissions for sneakers.txt*).
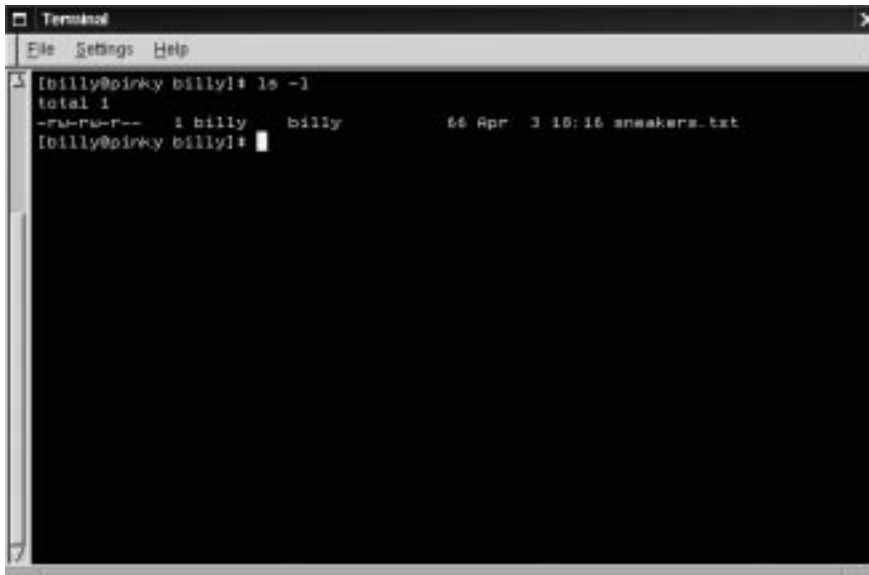
```
[billy@localhost billy]$ ls -l sneakers.txt
-rw-rw-r--    1 billy billy     150 Mar 19 08:08 sneakers.txt
```

There's quite a bit of detail here. We can see who can read (r) and write to (w) the file, as well as who created the file (billy) and to which group the owner belongs (billy).

Figure 3–16   Permissions for `sneakers.txt`

Other information to the right of the group includes the file name, date and time of
its creation as well as size.

How do all those dashes and letters fit together? It's not as hard to read as it might
seem. Let's take a look:

```
-rw-rw-r--
```

There are 10 slots in this column. The first slot represents the type of file. The re-
maining nine slots are actually three sets of permissions for three different categories
of users.

Those three sets are: the owner of the file, the group in which the file belongs and "others," meaning users and groups other than the owner of the file (billy) and those in billy's group (which is also billy).

Let's stretch out these file settings a bit:

```
   -     (-rw)   (-rw)  (r--)     1 billy billy
   |       |       |      |
  type   owner   group  others
```

The first item, which specifies the file type, can show one of the following:

- d -- a directory
- – -- a regular file (rather than directory or link)
- l -- a symbolic link to another program or file elsewhere on the system

Beyond the first item, in the following three sets, we'll see one of the following:

- r -- file can be read
- w -- file can be written to
- x -- file can be executed (if it's a program)

When we see a dash in owner, group or others, it means that particular permission hasn't been granted.

Let's look again at first column of sneakers.txt and identify its permissions. (See Figure 3–17, *A Closer View of Permissions*)

```
[billy@localhost billy]$ ls -l sneakers.txt
-rw-rw-r--    1 billy billy     150 Mar 19 08:08 sneakers.txt
[billy@localhost billy]$
```

Figure 3–17   A Closer View of Permissions

The file's owner, billy, has permission to read and write to the file; it's not a program, so billy doesn't have permission to execute it. The group, billy, has permission to read

and write to `sneakers.txt`, as well. Similar to the program notation for owner `billy`, there's no execute permission for group billy.

In the last set, we can see that those who aren't either the user billy or in the group called `billy` can read the file, but can't write to it or execute it.

We can use the `chmod` command to change a file's permissions.

Let's work a bit more on `sneakers.txt` to change the permissions with the `chmod` command.

The original file looks like this, with its initial permissions settings:

```
-rw-rw-r--    1 billy billy     150 Mar 19 08:08 sneakers.txt
```

As long as we're the owner of the file -- or we're logged into the root account -- we can change permissions in any combination of settings for the owner, group and others.

Right now, the owner (that's us) and our group (which is billy) can read and write to the file.

Anyone outside of our group -- for example can only read the file (`r--`).

Tip

Remember that file permissions are a security feature. Whenever you allow everyone to read, write to and execute files, you may be increasing your risk of tampering. As a rule, then, you should shy away from allowing everyone to read and write to a file.

In this case, however, let's say that we want to allow everyone to write to the file, so they can read it, write notes in it and save it. That means we'll have to change the change the "others" section of the file permissions.

Since we're the owner of the file, we don't have to `su` to root to do it. Let's take a look at the file first. At the shell prompt, type:

```
ls -l sneakers.txt
```

which gives us this file information:

```
-rw-rw-r--    1 billy billy     150 Mar 19 08:08 sneakers.txt
```

Now, type the following:
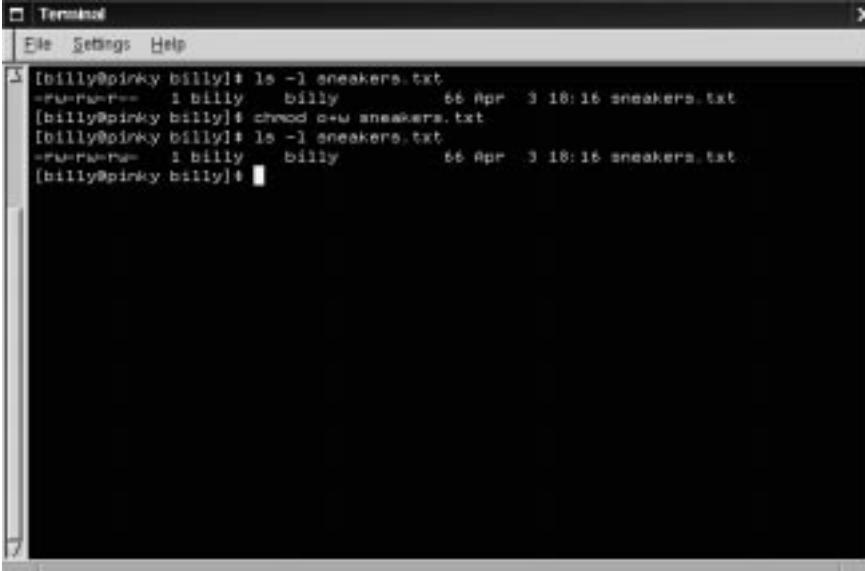
```
chmod o+w sneakers.txt
```

To check our results, we can list the file's details again. Now, the file looks like this:

```
-rw-rw-rw-    1 billy billy     150 Mar 19 08:08 sneakers.txt
```

There's our result: Now, everyone can read and write to the file (Figure Figure 3–18, *Changing Permissions for* sneakers.txt).

Figure 3–18   Changing Permissions for `sneakers.txt`

When we typed o+w, we were saying, "for others, add write permission to the file sneakers.txt."

If we want to remove all access permission from `sneakers.txt` (even though it's only a sketchy shopping list), we could use the `chmod` command to take away both the read and write permissions like so:

```
chmod go-rw sneakers.txt
```

and the result will look like this:

```
-rw-------    1 billy billy    150 Mar 19 08:08 sneakers.txt
```

By typing `go-rw`, then, we were saying "for the group and others, remove read and write permission to the file `sneakers.txt`."

You might think of these settings as a kind of shorthand when you want to change permissions with `chmod`, because all you really have to do is remember a few symbols and letters with the `chmod` command.

Here a list of what the shorthand represents:

3.12 Identities

`u` -- the user who owns the file (that is, the owner)

`g` -- the group to which the user belongs

`o` -- others (not the owner or the owner's group)

`a` -- everyone (`u`, `g`, and `o`)

3.12 Permissions

`r` -- read access

`w` -- write access

`x` -- execute access

3.12 Actions

`+` -- adds the permission

`-` -- removes the permission

`=` -- makes it the only permission

Want to test it out? Let's remove all permission from `sneakers.txt` -- for every-
one.
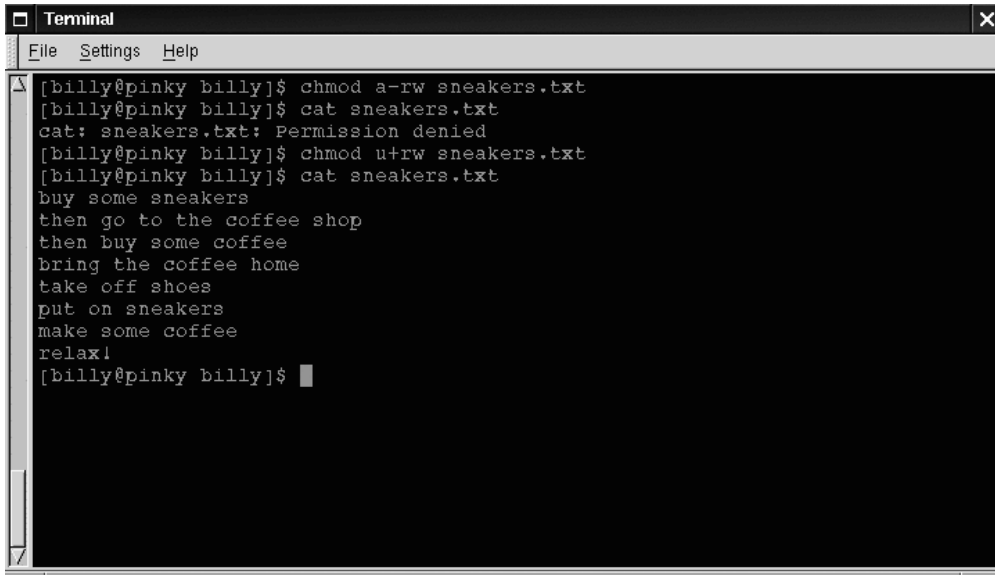
```
chmod a-rw sneakers.txt
```

Now, let's see if we can read the file:

```
[billy@localhost billy]$ cat sneakers.txt
cat: sneakers.txt: Permission denied
[billy@localhost billy]$
```

Guess it worked; even we can't get into the file. But since the file belongs to us, we
can always change permission to allow us read and write access. (See Figure 3–19,
Removing and Restoring Permissions)

```
[billy@localhost billy]$ chmod u+rw sneakers.txt
[billy@localhost billy]$ cat sneakers.txt
buy some sneakers
then go to the coffee shop
then buy some coffee
bring the coffee home
take off shoes
put on sneakers
make some coffee
relax!
[billy@localhost billy]$
```

Figure 3–19   Removing and Restoring Permissions

```
[billy@pinky billy]$ chmod a-rw sneakers.txt
[billy@pinky billy]$ cat sneakers.txt
cat: sneakers.txt: Permission denied
[billy@pinky billy]$ chmod u+rw sneakers.txt
[billy@pinky billy]$ cat sneakers.txt
buy some sneakers
then go to the coffee shop
then buy some coffee
bring the coffee home
take off shoes
put on sneakers
make some coffee
relax!
[billy@pinky billy]$
```

Here are some common examples of settings that can be used with chmod:

- g+w -- adds write access for the group

- o-rwx -- removes all permissions for others

- u+x -- allows the file owner to execute the file

- a+rw -- allows everyone to read and write to the file

- ug+r -- allows the owner and group to read the file

- g=rx -- lets the group only read and execute (not write)

By adding the -R option, we can change permissions for entire directory trees.

There's a slight twist, however, because we can't really "execute" a directory as we would an application. Instead, when we add or remove execute permission for a directory, we're really allowing (or denying) permission to search through that directory.

To allow everyone read and write access to the `tigger` directory in our login directory, we just type:

```
chmod -R a+rw tigger
```

But… If we don't allow others to have execute permission to tigger, it doesn't matter who has read or write access, because no one will be able to get into the directory -- unless they know the exact filename they want.

For example, let's type:

```
chmod a-x tigger
```

to remove execute access to all.

Here's what happens now when we try to `cd` to into `tigger`:

```
[billy@localhost billy]$ cd tigger
bash: tigger: Permission denied
[billy@localhost billy]$
```

Let's restore ours and our group's access.

```
chmod ug+x tigger
```

Now, if we check our work with `ls -dl` we'll see that only others will be denied access to tigger.

3.13 Fun with Numbers in `chmod`

Remember when we made a reference to the "shorthand" method of chmod? Here's another way to change permissions; it may seem a little complex at first - especially if math isn't your strong suit.

Let's go back to the original permissions for `sneakers.txt`.

```
-rw-rw-r--    1 billy billy     150 Mar 19 08:08 sneakers.txt
```

Each permission setting can be represented by a numerical value:

- r = 4
- w = 2
- x = 1
- - = 0

When these values are added together, the total is used to set specific permissions -
more specific than changing permissions with the alphabetical "shorthand."

In `sneakers.txt`, then, here are the numerical permissions settings:

```
-   (rw-)    (rw-)   (r--)
     |        |       |
   4+2+0    4+2+0   4+0+0
```

The total for the user is six, the total for the group is six and the total for others is
four. The permissions setting, then, is read as `664`.

If we want to change `sneakers.txt` so those in our group didn't have write ac-
cess, but could still read the file (as shown in Figure 3–20, *Removing Group Write
Permissions*), we'll have to remove the access by subtracting 2 from that set of num-
bers.

The numerical values, then, would become six, four and four -- or 644.
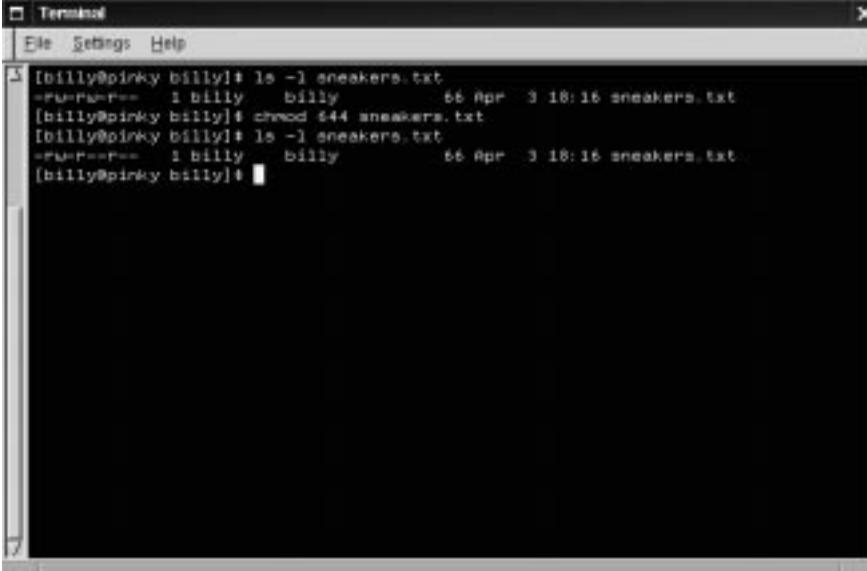
So we can type:

```
chmod 644 sneakers.txt
```

Let's check our changes by listing the file (`ls -l sneakers.txt`):

```
-rw-r--r--    1 billy billy     150 Mar 19 08:08 sneakers.txt
```

Figure 3–20   Removing Group Write Permissions

And there it is; now, neither the group nor others have write permission to sneak-ers.txt. To return the group's write access for the file, we can just add the value of w (2) to the second set of permissions.

```
chmod 664 sneakers.txt
```

Tip

Beware 666 and 777. Biblical implications aside, either of these settings will allow everyone to read and write to a file or directory. Such settings as these could allow tampering with sensitive files, so in general, it's not a good idea to allow these settings.

Here's a list of some common settings, numerical values and their meanings:

- `-rw-------` (600) -- Only the user has read and write permissions.

- `-rw-r--r--` (644) -- Only user has read and write permissions; the group and others can read only.

- `-rwx------` (700) -- Only the user has read, write and execute permissions.

- `-rwxr-xr-x` (755) -- The user has read, write and execute permissions; the group and others can only read and execute.

- `-rwx--x--x` (711) -- The user has read, write and execute permissions; the group and others can only execute.

- `-rw-rw-rw-` (666) -- Everyone can read and write to the file. Bad idea.

- `-rwxrwxrwx` (777) -- Everyone can read, write and execute. Another bad idea.

Here are a couple common settings for directories:

- `drwx------` (700) -- Only the user can read, write in this directory.

- `drwxr-xr-x` (755) -- Everyone can read the directory, but its contents can only be changed by the user.

Summary

You can change permissions with the `chmod` command by using letters or numbers. Type `chmod` *permissions* `file` to change permissions of a file or directory.

You've already come quite a distance in learning about your Red Hat Linux system -- from navigation to setting and changing permissions. Now, it's time to learn a little more about managing what you have on your system.

4   Managing Files and Directories

If you're a Linux newbie -- and there are many -- you may feel a little disoriented when you want to accomplish your first tasks.

Relax. If you've had any experience with other operating systems, learning Linux is a bit like learning to drive in a new country: Many of the ideas are the same, but some of the particulars are a bit different.

We'll go over several of these "rules of the road" in this chapter.

But there's one component of your new operating system you just can't do without: the **shell**. We've made numerous references to the shell -- as in "shell prompt," or "bash."

Now, it's time to learn a little more about this indispensable tool. But first, some background information…

4.1 Shell Collecting

In the olden days (we're talking '60s here), when AT&T's Dennis Ritchie and Ken Thompson were designing UNIX™, they wanted to create a way that humans could interact with their new system.

Operating systems at the time did come with **command interpreters**, which could take commands from the user and interpret them so the machine could use them.

But Ritchie and Thompson wanted something more, something which could offer better features than the command interpreters of the day.

Enter the **Bourne shell** (known simply as **sh**), created by S.R. Bourne, which fulfilled the goals of UNIX's creators.

Since the creation of the Bourne shell, other shells have been developed, such as the **C shell** (**csh**) and the **Korn shell** (**ksh**).

When the Free Software Foundation sought a royalty-free shell, developers began to work on the language behind the Bourne shell as well as some of the popular features from other shells available at the time.

The result was the **Bourne Again Shell** -- or **bash**.

By now, of course, you've probably seen the word `bash` when you've mistyped commands at the shell prompt (as in `bash:   somecommand:   command not found`).

In Chapter 3, *You Are Here*, when we talked about redirection and piping, we were also demonstrating the power of `bash`.

Tip

You can learn more about `bash` by reading the `bash` man page. At the shell prompt, type `man bash` (or you can save the file as a text file by typing `man bash | col -b > bash.txt`, which you can then open to read with an editor like `pico` or a pager like `less`. You can also print the file with `man bash | col -b | lpr`, but be warned: it's a large file. If you want more information, O'Reilly & Associates publishes *Learning the bash Shell*, by Cameron Newham and Bill Rosenblatt.

Although your system came with several different shells, `bash` is the default shell for Red Hat Linux.

You might think of `bash` as a fleet-footed office assistant who has made a habit of keeping notes on ways to fulfill commands quickly. This assistant also keeps "pointers" on how you like to customize the way you work.

These pointers `bash` keeps are referred to as **environment variables**.

The shell uses an "environment" in the same way we use an environment, like a kitchen. We work in our kitchen, arrange pots, pans, and spices. We know where the dishes are, how things operate.

The same can be said for `bash` and its environment. There's a basic arrangement to `bash` as there would be to just about any kitchen. For example, we'd expect to see pots in a kitchen in the same way that we would expect to see certain commands in `bash`.

That's the idea behind environment variables.

As long as your assistant has the right pointers, he'll fulfill your commands quickly.

Let's take a look at our environment variables. At the shell prompt, type:

```
env
```

Quite a few "shortcuts" `bash` uses, aren't there?

Each one of these helps `bash` customize the environment for you.

Among the most important environment variables is the **PATH** environment variable -- which defines what is known as the **default path**.

The **PATH** environment variable for our account `billy` might look something like this:

```
PATH=/usr/local/bin:/usr/X11R6/bin:/usr/bin:/bin:/usr/X11R6/bin:/home/billy/bin
```

It looks crowded, but the **PATH** statement is a great signpost, which points to where programs can be found.

Tip

Remember the earlier reference to the Filesystem Hierarchy Standard (see Section 3.4, *A Larger Picture of the Filesystem*)? The **PATH** statement is set according to that standard, and programs are installed in directories in accordance with the FHS as well. The result is that the **PATH** statement will enable `bash` to automatically find nearly any program, assuming it has been installed in accordance with the FHS.

4.2 Locating Files and Directories

There will be times when we know a file or directory exists but we won't know where to find it. Searching for a file or directory can be easier with the `locate` command.

With `locate`, we'll see every related file or directory which matches our search criterion. Let's say we want to search for all files related to the `finger` command.

```
locate finger
```

The `locate` command uses a database to check for files and directories which match the string `finger`.

Tip

To learn more about `locate`, read the `locate` man page
(type `man locate` at a shell prompt).

It's a handy command which works very quickly -- as long as the database is up to date. That database is automatically updated on a nightly basis, from `cron`. What's `cron`? It's a small program that runs in the background, performing various tasks -- such as updating the `locate` database -- at regularly scheduled intervals.

Tip

Tip: `cron` is a **daemon**. Daemons handle tasks in the back-ground. To read the `cron` man page, type `man cron` at the shell prompt.

So what happens if we:

•   Have more than one operating system on our machine, and switch between them -- causing us to halt and restart our Red Hat Linux system;

•   Shutdown and turn off our machine at the end of the day.

This might mean that `cron` rarely has a chance to update the **slocate** database, which is used to catalogue locations of files. But we can just update the database manually. Let's give it a try.

First, `su` to root (type **su** in at the prompt, then type your rootpassword).

Now, at the shell prompt, type:

```
updatedb
```

After perhaps a few minutes, the slocate database will be current.

4.3 Command History and Tab Completion

It doesn't take long before the thought of typing the same command over and over becomes unappealing, at best.

Linux users don't feel any differently about that, either. In Linux, since you can string together commands at the shell prompt, one minor typo in a couple lines of a command could mean that all that typing was in vain.

So there's a solution: It's called **command-line history**. By scrolling with the up and down arrow keys, we can find plenty of our previously typed commands -- including the ones with typos.

Let's try it by taking a look again at `sneakers.txt`. The first time, however, at the shell prompt, we'll type:

```
cat sneakrs.txt
```

Nothing happens, of course, because there is no `sneakrs.txt` file. No problem. We'll just use the up-arrow key to bring back the command, then use the left-arrow key to get to the point where we missed the "e." Insert the letter and press [Enter] again.

Voila! We now see the contents of `sneakers.txt`.

By default up to 500 commands can be stored in the bash command-line history file.

Tip

By typing the env command at a shell prompt, we can see the environment variable that controls the size of the command-line history. The line which reads, HISTFILE-SIZE=500 shows the number of commands that bash will store.

The command-line history is actually kept in a file, called .bash_history in our login directory. We can read it in a number of ways: by using pico, cat, less, more, and others.

Be prepared, though: the file can be pretty long.

Let's read it with more:

```
more .bash_history
```

To move forward a screen, press [Space]; to move back a screen, press [B]; to quit, press [Q].

Tip

Want to find a command in your history file without having to keep hitting the arrow keys or page through the history file? Use `grep`, a powerful search utility. Here's how you can quickly find a previously used command: Let's say you're searching for the command that was something like `cat sneak`-something. You've used the command and think it might be in your history file. At the shell prompt, then, type

```
history | grep sneak
```

In addition to the command you've just typed, you'll see your sought-after command, as well, because `grep` searched through your history file for every instance in which the word "sneak" appeared. You can learn more about `grep` later in this chapter, when we discuss tools which can help you read files.

Another time-saving tool is known as **command completion**. If you type part of a file, command or pathname then press the [Tab] key, `bash` will present you with either the remaining portion of a the file/path, or a beep. If you get a beep, just press [Tab] again to obtain a list of the files/paths that match what's been typed so far.

For example, if you forget the command `updatedb`, but remember a portion of the command, you can `su` to root, then at the shell prompt, type `up`, press the [Tab] key twice and you'll see a list of possible completions, including `updatedb` and `uptime`. By adding the letter "d" to `up` and pressing [Tab] again, your command is completed for you.

So even if the machine is turned off at the end of the day, it's not too hard to update the `slocate` database: The chances are good that the command will be stored in the command-line history or can be completed with command completion (as long as you remember the start of the pathname for the command).

4.4 Identifying and Working with File Types

If you're new to Linux, it won't take long before you begin seeing files with extensions that may seem foreign. A file's extension is the last part of a file's name, after the final dot (in the file `sneakers.txt`, "`txt`" was that file's extension).

Here's a brief listing of extensions and their meanings:

4.4.1 Compressed/Archived Files

- `.Z` -- a compressed file
- `.tar` -- an archive file (short for *tape archive*)
- `.gz` -- a compressed file (gzipped)
- `.tgz` -- a tarred and gzipped file

4.4.2 File Formats

- `.txt` -- a plain ASCII text file
- `.html/.htm` -- an HTML file
- `.ps` -- a PostScript file; formatted for printing
- `.au` -- an audio file
- `.wav` -- an audio file
- `.xpm` -- an image file
- `.jpg` -- a graphical or image file, such as a photo or artwork
- `.gif` -- a graphical or image file
- `.png` -- a graphical or image file

4.4.3 System Files

- `.rpm` -- a Red Hat Package Manager file
- `.conf` -- a configuration file
- `.a` -- an archive file

- `.lock` -- a "lock" file; determines whether a program is in use

4.4.4 Programming and Scripting Files

- `.h` -- a C and C++ program language header file
- `.c` -- a C program language source code file
- `.cpp` -- a C++ program language source code file
- `.o` -- a program object file
- `.pl` -- a Perl script
- `.tcl` -- a TCL script
- `.so` -- a library file

But file extensions are not always used, or used consistently. So what happens when a file doesn't have an extension, or the file doesn't seem to be what the extension says it's supposed to be?

That's when the `file` command can come in handy.

In Chapter 3, *You Are Here*, we created a file called `saturday` -- without an extension. Using the `file` command, we can tell what the file is by typing:

```
file saturday
```

and we'll see it's a text file. Any file that's designated a text file should be readable using `cat`, `more`, or `less`.

Tip

To learn more about `file`, read the `file` man page by typing `man file`.

And speaking of reading files…

There are plenty of ways to read files in Linux. In Chapter 3, *You Are Here*, for example, we covered the pagers `more` and `less` -- they're called pagers because

you can "page" through documents one screen at a time. We also learned how we can not only view but manipulate files with `cat`.

But there are even more options when it comes time to take a look at README files, man pages or documents you've created.

You have a number of tools to help you read text files, among them, the text editors pico, emacs, and vim, the pagers `more` and `less`, and the viewers `head`, `tail`, `cat`, and `grep`.

Let's take a look at some of the features in these tools.

The `less` Command

In Chapter 3, *You Are Here*, we introduced to the pager `less`. Less is the pager that's used to display man pages.

Let's view the man page for `less` to see `less` in action.

```
man less
```

To move forward a screen, press [Space]; to move back a screen, press [B], and to quit, press [Q].

There are other powerful features to `less`, as well, including the ability to scroll horizontally and specify the number of lines to scroll.

The `more` Command

Odd as it may seem, `more` offers less than `less` (actually, `less` was inspired by `more`).

Let's take a look at the man page for `more`, but this time, we'll open the page using `more` -- by piping `man`'s output to `more`.

```
man more | more
```

It may not look too different at first, but there are fewer enhancements to `more` than to `less`. Probably the most striking difference at first is the lack of a way to go backwards in a document -- although moving forward by pressing [Space] and quitting by pressing [Q] are the same.

The `head` Command

You can use the `head` command if you just want to look at the beginning of a file. The command is:

```
head <filename>
```

`Head` can be useful, but because it's limited to the first several lines, you won't know how long the file actually is. By default, you can only read the first 10 lines of a file, although we can specify the number to see more by typing:

```
head -20 <filename>
```

Read `head`'s man page (`man head`) for more information. You'll probably find that `less` or `more` are more helpful, because you can page through the file if you find that the information you're looking for is further into the file than you originally thought.

The `tail` Command

The reverse of `head` (obviously), is `tail`. With (`tail`), you can review the last 10 lines of a file.

The `cat` Command

The command `cat`, short for concatenation, will dump the contents of the entire file on the screen. Using `cat` can be handy if the file is fairly short, such as when we created `sneakers.txt`. But if a file is fairly long, it will easily scroll past you on the screen, since `cat` displays the whole file.

The `grep` Command

The `grep` command is pretty nifty for finding specific character strings in a file. Let's say we want to find every reference we made to "coffee" in the file `sneakers.txt`, which we created in our login directory. We could type:

```
grep coffee sneakers.txt
```

and we would see every line in which the word "coffee" could be found.

Tip

Unless otherwise specified, grep searches are **case sen-
sitive**.   That means that searching for *Coffee* is different
than searching for *coffee*. So among grep's options is -i,
which allows you to make a case-insensitive search through
a file. Read the grep man page for more about this com-
mand.

I/O Redirection and Pipes

And don't forget about using pipes and output redirection when you want to store
and/or print information to read at a later time.

You can, for example, use grep to search for particular contents of a file, then have
those results either saved as a file or sent to a printer.

To print the information about references to "coffee" in sneakers.txt, for exam-
ple, just type:

```
grep coffee sneakers.txt | lpr
```

This command behaves similar to the command ls -al /etc | more, which
you may have used in Chapter 3, *You Are Here* to list the contents of the /etc direc-
tory then send the results through the more command for viewing on the screen.

Tip

Remember the distinction of using > and >>: using > will
overwrite a file, while >> appends the information to a file.
Usually, unless you're certain you want to, it's safer to use
>>, because you won't lose potentially valuable information
(though you may have to edit the file if you didn't want to
append information to it).

Wildcards and Regular Expressions

What if you forget the name of the file you're looking for? You can't say to your computer, "Find a file called 'sneak' or 'sneak-something'."

Well, yes you can, in a way. Using **wildcards** or **regular expressions**, you can perform actions on a file or files without knowing the complete filename. Just fill out what you know, then substitute the remainder with a wildcard.

Tip

To read more about wildcards and regular expressions, take a look at the bash man page (`man bash`). Remember that you can save the file to a text file by typing `man bash | col -b > bash.txt`. Then, you can open and read the file with `less` or `pico` (`pico bash.txt`). If you want to print the file, be prepared: It's quite long.

We know the file's called "sneak-something.txt," so just type:

```
ls sneak*.txt
```

and there's the name of the file:

```
sneakers.txt
```

You'll probably use the asterisk (`*`) most frequently when you're searching. The asterisk will search out everything that matches the pattern you're looking for. So even by typing:

```
ls *.txt
```

or:

```
ls sn*
```

you'd find `sneakers.txt` -- except that as time goes on, there will be more text files, and they'll all show up because they match the pattern you're searching for.

It helps, then, to narrow your search as much as possible.

One way to narrow that search might be to use the question mark symbol (`?`). Like the asterisk, using `?` can help locate a file matching a search pattern.

In this case, though, `?` is useful for matching a single character -- so if you were searching for `sneaker?.txt`, you'd get `sneakers.txt` as a result -- and/or `sneakerz.txt`, if there were such a filename.

When an asterisk, for example, just happens to be part of a filename, such as might be the case if the file `sneakers.txt` was called `sneak*.txt`, that's when regular expressions can come in handy.

Regular expressions are more complex than the straightforward asterisk or question mark.

Using the backslash (\), you can specify that you don't want to search out *everything* by using the asterisk, but you're instead looking for a file with an asterisk in the name.

If the file is called `sneak*.txt`, then, type:

```
sneak\*.txt
```

Here is a brief list of wildcards and regular expressions:

- `*` -- Matches all characters
- `?` -- Matches one character in a string (such as `sneaker?.txt`)
- `\*` -- Matches the * character
- `\?` -- Matches the ? character
- `\)` -- Matches the) character

You can also use wildcards for more than searching: they can come in handy when you want to move and rename files. And regular expressions can help you rename files with characters like * and ? in them.

For more on that, read on.

4.5 Copying, Moving and Renaming Files and Directories

By now, you've learned a little about the structure of the filesystem; and you've learned how to create files and directories.

But just because you know how to create files and directories doesn't mean that you're stuck with the changes you've made. What if you want to rename and/or move files and directories?

Let's start with the copy command.

4.5.1 Copying Files

Like so many Linux features, you have a variety of options from which to choose when you want to manipulate files and directories. You can also use wildcards when you're copying, moving, or deleting files and directories.

Basically, the copy command is not much more complex than typing:

```
cp <source> <destination>
```

so to copy the file `sneakers.txt` to the directory `tigger` in your login directory, just type:

```
cp sneakers.txt tigger
```

Notice that you also used relative pathnames to copy the file. You can use both relative and absolute pathnames with `cp`. Our login directory is the parent of the directory `tigger`; meaning that `tigger` is one directory down from ours.

Read the `cp` man page (`man cp`) for a full list of the options available with `cp`. But among the options you can use with `cp` are:

- `-i` -- interactive. Prompts you to confirm if the file is going to overwrite a file in your destination. This is a handy option because it can help prevent you from making mistakes.

- -r -- recursive. Rather than just copying all the files and directories, copies the whole directory tree, subdirectories and all, to another location.

- -f -- force. Copies without prompting you for confirmation that the file should be overwritten. Unless you're sure you want to force the copy, you probably don't want to make friends with this option right now.

- -v - verbose. Will show the progress of the files being copied.

Just by using cp alone, you won't see much when the command is executed. Using an option, such as -i, can make the process a little more useful, because if you want to copy a file to a location that already has a file with the same name, you'll be asked first if you really want to overwrite -- meaning replace -- the file that's already there.

Tip

Remember that among your options is -f (force), which can overwrite files without asking you if you're certain. Make sure, when you use the force option, that you *really* want to overwrite a file.

Now that we have the file sneakers.txt in the tigger directory, let's use cp -i to copy the file again to the same location.

```
[billy@localhost billy]$ cp -i sneakers.txt tigger
cp: overwrite 'tigger/sneakers.txt'?
```

To overwrite the file that's already there, press [Y] and then [Enter]. Don't want to overwrite the file? Now's the time to press [N] and [Enter].

4.5.2 Moving Files

To move files, use the mv command (man mv), which is similar to the cp command, except that with mv the file is physically moved from one place to another, instead of being duplicated, as with cp.

Common options available with mv include:

- `-i` -- interactive. Will prompt you if the file you've selected will overwrite an existing file in the destination directory. This is a good option, because like the `-i` option in `cp`, you'll be given the chance to make sure you want to replace an existing file.

- `-f` -- force. Overrides the interactive mode and moves without prompting. Unless you know what you're doing, this option doesn't play nice; be very careful about using it until you become more comfortable with your system.

- `-v` -- verbose. Shows a list of the files being moved.

If you want to move a file out of your home directory and into another directory, you would type:

```
mv sneakers.txt tigger
```

or, `mv sneakers.txt /home/billy /home/billy/tigger` using absolute pathnames.

4.5.3 Renaming Files

Actually, we've already covered half of renaming, because when you copy or move files, you can also rename.

To copy the file `sneakers.txt` from our login directory to our `tigger` subdirectory, just type:

```
cp sneakers.txt tigger
```

To copy and rename that file from `sneakers.txt` to `piglet.txt`, type:

```
cp sneakers.txt tigger/piglet.txt
```

To *move* and rename the file, just substitute `mv` for `cp` in the above example.

If you `cd` to `tigger` and use `ls`, you'll see the file `piglet.txt`.

If you just want to rename the file and keep its location, just `mv` in your current directory:

```
mv  sneakers.txt piglet.txt
```

4.5.4 Deleting Files and Directories

We talked about creating files with the touch command and by using redirection in Chapter 3, *You Are Here*. And we created the directory tigger using mkdir.

But we haven't discussed how to delete files and directories.

Deleting files and directories with the rm command (man rm) is a straightforward process.

Let's take our new file piglet.txt, and delete it from the tigger directory with the rm command:

```
rm piglet.txt
```

What happens if we didn't really want to get rid of it? Too late! Again, that's where the -i (interactive) option comes in handy, because it gives a second chance to think about whether we really want to toss the file.

```
[billy@localhost billy]$ rm -i piglet.txt
rm: remove 'piglet.txt'?
```

You can also delete files using the wildcard *, but be careful, because you can easily delete files you didn't intend to throw away.

To remove a file using a wildcard, you would type:

```
rm pig*
```

You can also remove more than one file in one command, as in:

```
rm piglet.txt sneakers.txt
```

Options for removing files -- and directories -- include:

- -i -- interactive. Prompts you to confirm the deletion. This is good.
- -f -- force. Overrides interactive mode and removes the file(s) without prompting. This might not be good, unless you know exactly what you're doing.

- -v -- verbose. Shows a listing of files as they're being removed.

- -r -- recursive. When removing directories, will remove all of the files and the subdirectories of the specified directory. This can also get rid of an empty directory.

To remove directories with rm, you must specify the -r option.

For example, if you want to recursively remove the directory tigger you would type:

```
rm -r tigger
```

And if you want to combine options, such as forcing a recursive deletion, you can type:

```
rm -rf tigger
```

<div style="text-align:center">

CAUTION

</div>

Be careful! rm is a powerful command, and can delete your entire system! If you're root and you type the simple command rm -rf / you're sunk -- like a snake eating its tail, the command will recursively remove everything on your system.

A safer alternative to using rm for removing directories is the rmdir command. With this command, you won't be allowed to use recursive deletions, so a directory which has files in it won't be deleted.

Read the rmdir man page by typing man rmdir to find out more about the command.

5 Getting Online

New to Red Hat Linux 6.1 is RP3, the Red Hat PPP Dialer utility which can make easy work of connecting to the Internet or to some other network devices. You can use RP3 to connect your modem to a PPP account, then view statistics about the connection, such as connect time, speed and cost of the current selection.

In many cases, your only requirement for establishing a connection is having a valid Internet account. RP3 can also be used to monitor and control other networking devices, such as ethernet cards.

5.1 Using RP3 to Connect

You can start RP3 by going to the GNOME **Main Menu Button** => **Panel** => **Add applet** => **Network** => **RH PPP Dialer**. The application will then dock into the **Panel**. Another way to start the application and dock it into the **Panel** is to right-click on the **Panel**, then go to **Add applet** => **Network** => **RH PPP Dialer**.

You can also launch RP3 -- without docking it into the **Panel** -- from either the **GNOME Panel** (by going to **Internet** => **RH PPP Dialer**) or from the shell prompt, by typing `/usr/bin/rp3` in an Xterm window.

To see all of your user-controllable network devices, such as a modem and ethernet card, from the **Main Menu Button** go to **Internet** => **RH Network Monitor**. Similar to the two previous methods, the application will appear on the desktop, rather than docking into the **Panel**.

The **Add New Internet Connection** dialog allows you to easily configure your system for an Internet connection, helping you to locate your modem, enter the account information and connect.

After you have successfully configured your Internet account -- or if you already have devices which are currently configured for network use -- you'll see a screen like Figure 5–1, *The* **Choose an Interface** *Dialog*, which will show your currently configured devices or accounts.

Figure 5–1   The Choose an Interface Dialog
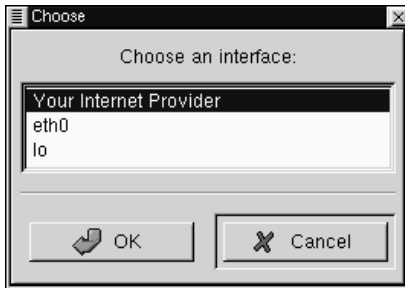
5.1.1 Creating a New Internet Account

To create a new PPP account, start RP3. If you don't already have devices configured, you'll see the **Add New Internet Connection** dialog (see Figure 5–2, *Creating a New Internet Account*).

Figure 5–2   Creating a New Internet Account

If there are devices detected, such as an ethernet card, you can simply choose a listed device so that the application is running on the **Panel** or the desktop, as in Figure 5–3, *The* RP3 *Monitor*. Once the application is running, right-click on the application to open the options menu and select **Configure PPP...**

Figure 5–3   The RP3 **Monitor**
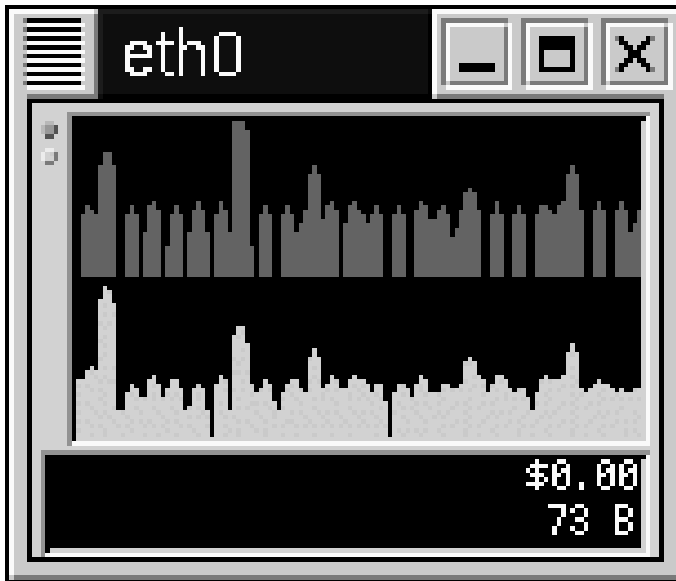
Next, type your root password when you're prompted and click the **OK** button to continue to the opening screen for the **Add New Internet Connection** tool.

To proceed, select the **Next** button; to quit the utility without making any changes or additions, use the **Cancel** button. You can cancel without saving changes at any time throughout the process; you can also return to a previously completed screen by using the **Back** button.

Your next task is to configure a modem. RP3 will try to automatically detect a modem on your system. Regardless of its success, you'll find a modem configuration dialog, similar to that in Figure 5–4, *Selecting a Modem Dialog*.

Figure 5–4 Selecting a Modem Dialog
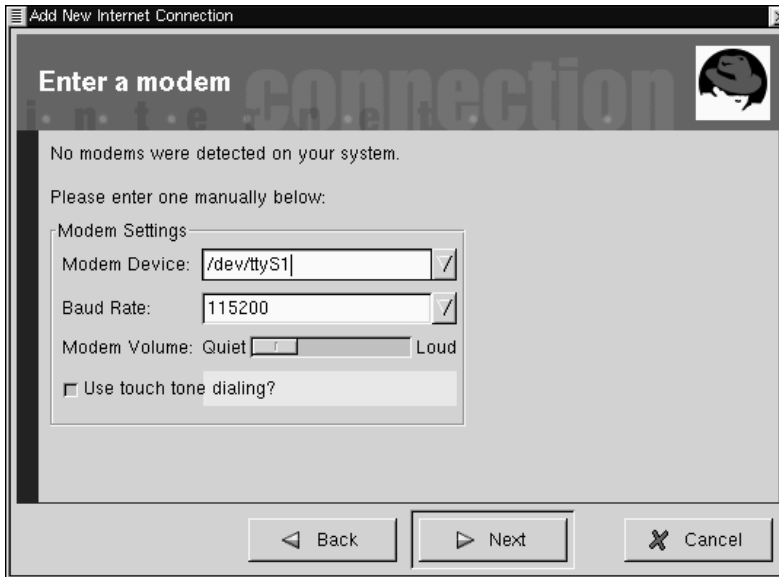
In the modem settings dialog, you can specify the location of your modem, its highest baud rate, your preference regarding the volume of the call and whether to use touch tone dialing.

Like the modemtool utility, you can specify the serial port to which your modem is connected, and a link will be created to /dev/modem. Although RP3 attempts to detect the correct port, you can adjust the results in this dialog.

If you know that your modem is attached to COM2 under Windows, for example, its location in Linux will be /dev/ttyS1. Here is a comparison of serial line names in MS-DOS (and Windows) and Linux:

- COM1 = /dev/ttyS0

- COM2 = /dev/ttyS1

- COM3 = /dev/ttyS2

- COM4 = /dev/ttyS3

Using the drop-down list in the **Modem Device** box, enter the serial port to which your modem is connected. Then, in **Baud Rate**, enter the fastest recommended speed for your modem (a typical 56k device, for example, would have a baud rate of 115200).

You can also select your volume and dialing preferences in this dialog. Dragging the bar to **Loud** in the **Modem Volume** entry will let you hear everything that goes on during a connection -- from dialtone to the connection negotiation. Note, however, that some modems will allow only a sound on/sound off configuration, rather than gradations of volume.
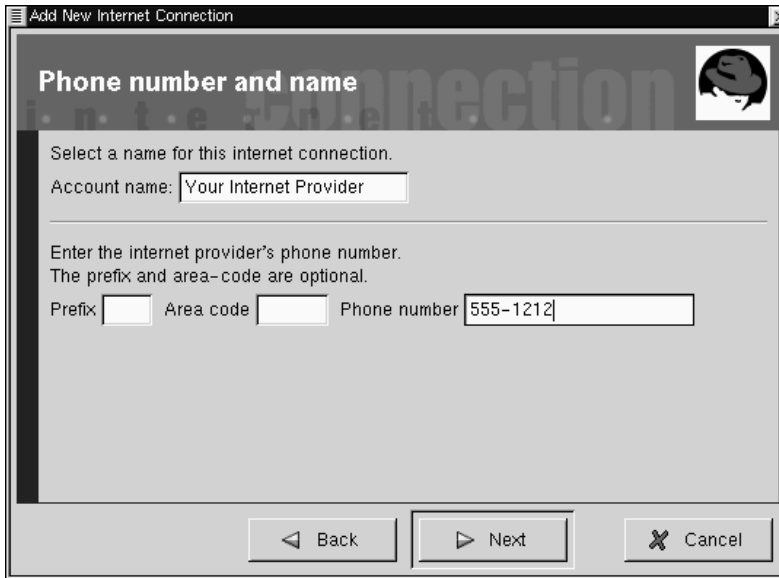
The **Use touch tone dialing** button, when selected, will dial using touch tone; when deselected, you'll dial using pulse.

Once you're satisfied with your settings, click on **Next** to proceed to the following screen.

In the **Phone number and name** dialog (see Figure 5–5, *Adding a Name and Phone Number*), you should enter both the name you would like to give your account and the telephone number of your Internet service provider. Ordinarily, it isn't necessary to provide the prefix or area code of your ISP.

If you must dial a "9" to reach an outside line, enter this number in the **Prefix** field. To disable call waiting, type ***70** in this field.

RP3 uses an application called Wvdial for some of its functions, such as dialing and authentication with the PPP server. To learn about other dialing options which might be available, type `man wvdial` in an Xterm window. (Use the [Space] bar to move forward through the document a "page" at a time; use the [B] key to move back, and use the [Q] key to quit.)

Figure 5–5   Adding a Name and Phone Number

The dialog which follows, **User name and password**, asks you to provide the name you use to sign on to your ISP, as well as your login password for the service. Your password will appear as a series of asterisks.

In the **Other Options** dialog, you can choose a pre-selected ISP or a **Generic Account** if your ISP isn't listed. If your ISP is listed, you should select the appropriate entry, due to variations in connection protocols which require specific network negotiations between your system and your ISP's.

You can check your information in the **Create the account** dialog which follows (see Figure 5–6, *Checking Your Entries*). Presented here will be the information you had provided in the previous screens.

Figure 5–6   Checking Your Entries
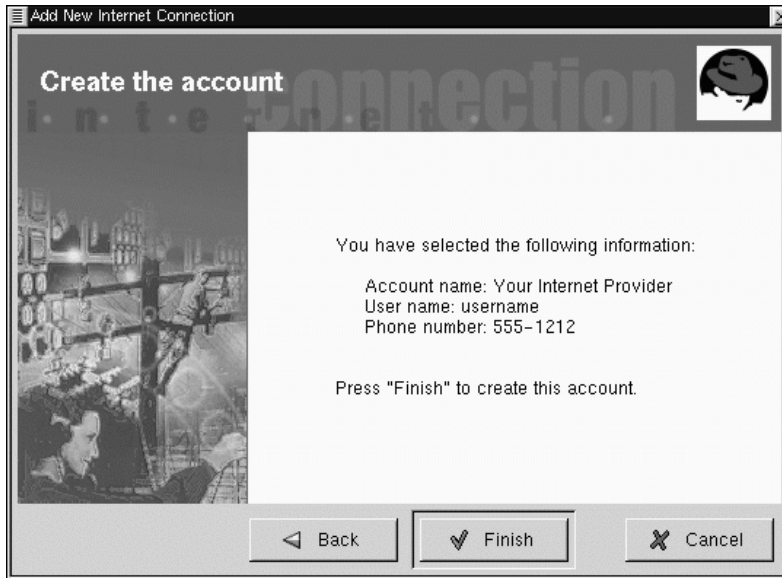
If you're satisfied with your entries, click **Finish** to create the account; if you want to modify the information, select the **Back** button to go to a previous screen.

The configuration tool modifies multiple files. The main files which are affected are `/etc/sysconfig/network-scripts/ifcfg-ppp*` and `/etc/wv-dial.conf`. If you find it necessary to modify an existing entry, however, you can edit that entry with RP3, rather than hand-edit a file. See Section 5.2, *Customizing and Configuring* for more information on customizing the settings.

5.1.2 Connecting and Disconnecting

Once you have configured your PPP account, start RP3 (right-click on the **Panel** => **Add applet** => **Network** => **RH PPP Dialer**).

When the application starts, you'll be asked to select which interface you want to activate; choose the name of the PPP service you created. Next, he application will start by showing you the monitor for the connection (see Figure 5–3, *The* RP3 *Monitor*).

To begin your online session, left-click on the monitor and, in the dialog which asks whether you want to start the interface, select **Yes**. (You can also right-click on the monitor, then drag your cursor to the **Connect to...** entry showing your ISP account name.)

Similar to the procedure for connecting, you can disconnect to end your online session by left-clicking on the monitor, then choosing **Yes** to stop the interface. (Also, you can use the **Disconnect from...** entry by right-clicking on the monitor.)

5.1.3 Working with the RP3 **Monitor**

When the RP3 monitor is docked on the **Panel**, it will appear similar to Figure 5–7, *The* RP3 *Monitor Docked (second icon from left)*.

Figure 5–7   The RP3 **Monitor Docked (second icon from left)**
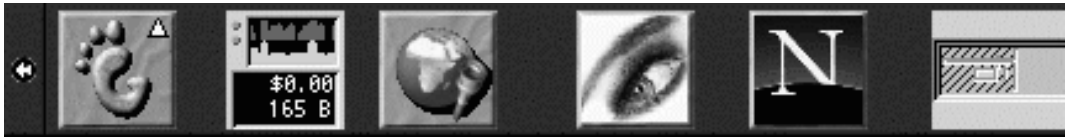
During the time you're connected, the RP3 monitor will display two bar graphs showing connection activity: the top graph will show data transmitted, and the bottom graph will show data received (see Figure 5–8, *Closeup of the* RP3 *Monitor*).

Figure 5–8   Closeup of the RP3 **Monitor**
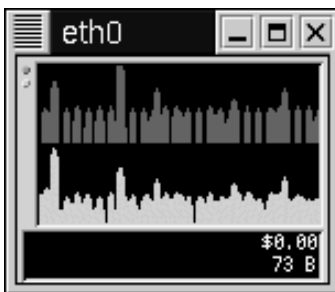
Beneath the graphs, you'll find two text entries which can display either your total connection time and the bytes received per second or your total cost per session, as well as the bytes received per second.

5.2 Customizing and Configuring

Once you've set up a PPP account, you can further configure and customize RP3.

Some settings which are displayed in RP3can be determined by your preferences. To view or change the settings, right-click on the application and, in the menu which appears, left-click once on **Properties**.

In the **Properties** dialog (see Figure 5–9, *The* **Properties** *Dialog*), you'll find the following entries:

Figure 5–9   The Properties Dialog
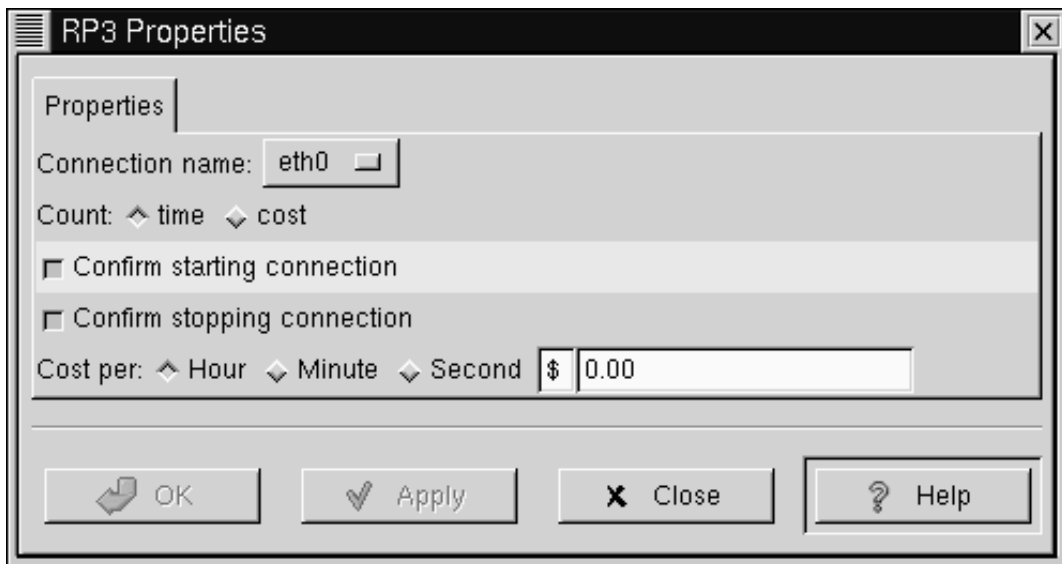
- **Connection name:** Allows you to switch between network connections that you have previously configured.
- **Count:** Allows you to specify whether you want to view your time online or the costs of your online session.
- **Confirm starting connection**: Confirms your intention to start a connection.
- **Confirm stopping connection**: Confirms your intention to stop a connection.

- **Cost per:** You can elect to view the costs of your session by the hours, minutes, or seconds. To view this feature, you must select the **Count: cost** option.

The settings you'll find in the **Properties** dialog relate to the configured devices and accounts. You can also customize and configure account information by right-clicking on the RP3 monitor and selecting the **Configure PPP** entry.

Figure 5–10   The Properties Dialog
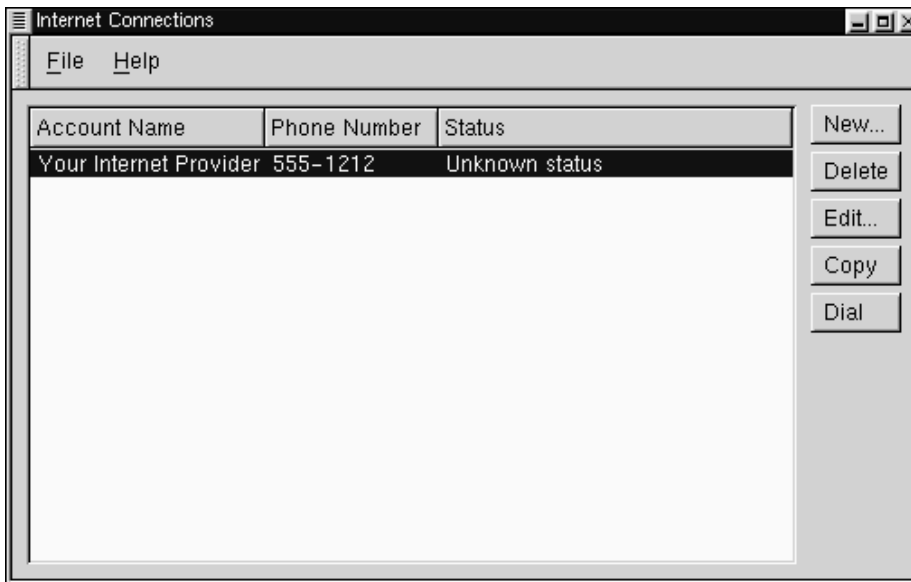
In the **Internet Connections** dialog which appears, you can add a new account, edit or delete a current account, copy the account to create a new entry, and manually dial the account.

Selecting the **New** entry will begin the Configure PPP tool, allowing you to enter new account information.

When you choose to **Edit** a current entry, you can change account details in the **Account Info** tab. Among the fields you can change are your login name, your password, and ISP telephone number.

In the **Modem** tab of the dialog, you can specify the location of your modem (such as /dev/ttyS1). You can either type in the information, or click on the **Find Modem...** button, which will attempt to automatically detect your modem.

6   Using Gnome-RPM and RPM

6.1 The benefits of RPM

RPM stands for *R*ed *H*at Package *M*anager. At its heart, RPM is one of the most convenient technologies available for Red Hat Linux. That convenience extends beyond the user; Red Hat encourages other vendors to consider the ease of use behind the RPM system for their own products (the system may be distributed under terms of the GPL).

Want to install or remove an application? No problem. Upgrade an application you already have? It's easy. By typing a few simple commands or pressing a few buttons, the process will be taken care of for you.

RPM also features useful tools to view the contents of a **package**, as well as other essential file and system maintenance features.

From its earliest days, RPM was intended to be an easy and secure way to get packages on an off a user's system. Often, the only requirement to installing an RPM package -- sometimes referred to as a **pre-compiled binary** if it's an application -- is that you're root and that you press the right buttons!

You can choose to work with RPM packages from either a graphical front-end or the shell prompt. Gnome-RPM, which was introduced in Red Hat Linux 6.0, is a graphical RPM tool. An alternative to Gnome-RPM are various RPM commands you can issue from the shell prompt in an Xterm window.

This chapter will cover some of the basic features of RPM use from Gnome-RPM and from the shell prompt. For more in-depth information, about both Gnome-RPM and RPM technology, refer to the related chapters in the *Official Red Hat Linux Reference Guide.*

Please Note

Because RPM makes changes to your system, for most func-
tions you must be root in order to use RPM from either the
shell prompt or through Gnome-RPM.

6.1.1 Changing to Root

Changing to root can be easily accomplished from within your user account.

You can temporarily switch from your user account to root by using the su com-
mand. Once you're root, you can install, delete and work with packages and perform
maintenance on system files.

Remember, however, that when you're root you can inadvertently damage your sys-
tem, so use caution!

Just type su at the shell prompt. You'll see something like this:

```
[you@localhost you]$su
Password: yourrootpassword
[root@localhost you]#
```

Now that you're root, you can start Gnome-RPM or perform your RPM command from
the shell prompt.

6.2 Using Gnome-RPM

Gnome-RPM was written by James Henstridge (james@daa.com.au); RPM 3.0
support was written by Red Hat and additional rpmfind code was written by Daniel
Veillard. Gnome-RPM is a fast, convenient tool for working with RPMs; with such fea-
tures as **Web find**, it's also a powerful utility which allows you to easily search for
new RPMs on the Internet.

You can start Gnome-RPM (which is also referred to as gnorpm) from the shell prompt in your user account, and from both the shell prompt and the GNOME Panel if you have logged in as root (rather than performing an su from your user account).

To start Gnome-RPM from the Panel, go to the **Main Menu Button** => **System** => **Gnome-RPM**.

To start Gnome-RPM from the shell prompt, type gnorpm. When the application starts, you'll see a window similar to Figure 6–1, *Main* Gnome-RPM *Window*.

Figure 6–1   Main Gnome-RPM **Window**
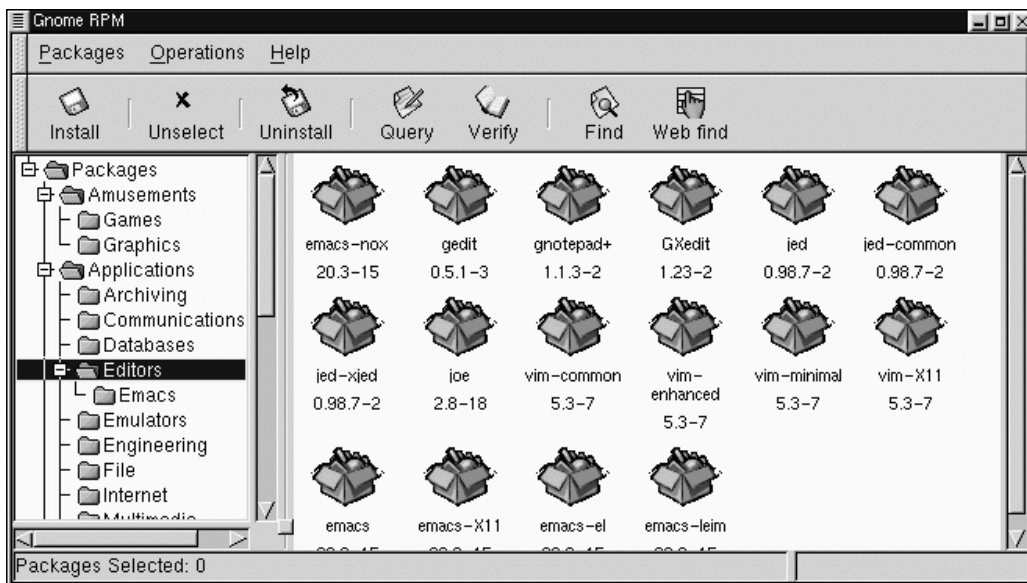
6.2.1 The Package Window

The main window will appear similar to Figure 6–1, *Main* Gnome-RPM *Window*.

- Package Panel - On the left; allows you to browse and select packages on your system.

- Display window - To the right of the package panel; shows you contents from folders in the panel.

- Toolbar - Above the display and panel; a graphical display of package tools.

- Menu - Above the toolbar; contains text-based commands, as well as help info, preferences and other settings.

- Status bar - Beneath the panel and display windows; shows the total number of selected packages.

6.2.2 The Package Display

Each folder icon in the tree view at left represents a group of packages. Each group can contain subgroups. Groups are used to place packages that perform similar functions in similar locations.

This tree view is expandable. A folder which appears with a **+** next to it indicates that there are subfolders within that category.

To view the packages and subgroups within a group, click once on a folder or a **+** with your left mouse button. The display window will then show you the contents of that folder.

By default, you'll see icons to represent the packages. You can change customize your view by selecting **View as list** from the **Interface** tab you'll find under **Operations** => **Preferences**.

6.2.3 Selecting Packages

To select a single package, click on it with the left mouse button. You'll notice that highlighting will appear around the package's title (as shown in Figure 6–2, *Selecting Packages in* Gnome-RPM). To unselect it, either click on an empty space in the display panel with the left mouse button, or click on the Unselect button on the toolbar.

Figure 6–2   Selecting Packages in Gnome-RPM
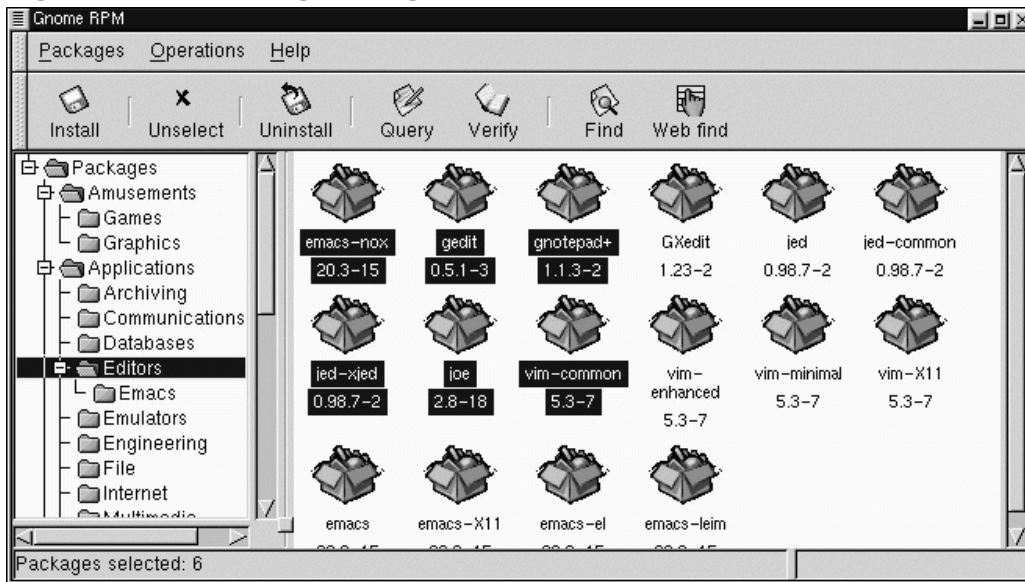
To select more than one package incrementally, left-click with your mouse button, while holding down the [Ctrl] key; you'll see highlighting around each additional selection.

The status bar at the bottom of Gnome-RPM will display the total number of packages you have selected.

6.3 Installing New Packages

You can install new packages by choosing the **Install** button from the toolbar.

Gnome-RPM provides you with a choice of package views (as shown in Figure 6–3, *The Install Window*). You can **filter** your views for

- All packages

- All but installed packages
- Only uninstalled packages
- Only newer packages
- Uninstalled or newer packages

You can switch the display of packages by using the drop-down bar at the top of the window.

Figure 6–3   The Install Window
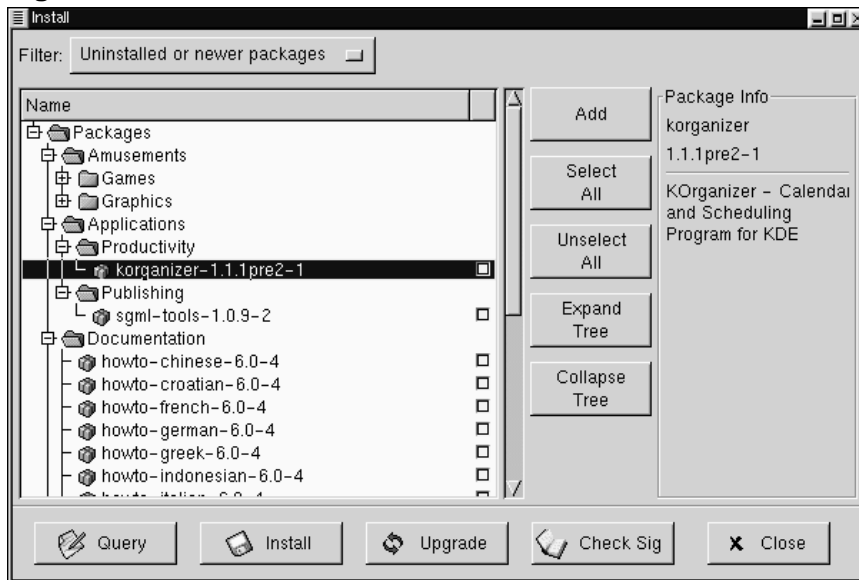
Choose the **Add** button. By default, if your CD-ROM is mounted with a Red Hat Linux CD, Gnome-RPM will search in /mnt/cdrom/RedHat/RPMS for new packages. (You can find this default option in the **Install Window** tab of the **Preferences** dialog. See Section 6.4, *Basic Configuration* for more information on this feature.)

If no packages are available in the default path, you'll be presented with an **Add Packages** window from which you can select the appropriate location of your new package.

To choose a package, double-click on it with your mouse, or add a check in the checkbox next to the package, then click on the **Add** button.

Within this **Install** window, you have the option of installing or upgrading packages: If the package has not been installed on your system, you'll find an **Install** button at the bottom of the window.

If the package already exists on your system and you're querying a newer release, the **Package Info** window provides an **Upgrade** button, which will perform an upgrade to newer releases.

You can also "drag and drop" packages from the GNOME File Manager and place them into the **Install** window. To accomplish this, open the File Manager (**GNOME Panel** => **File Manager**). Locate the directory in which your packages can be found, left-click once on the package name and "drag" the file by keeping your finger pressed on the mouse button, releasing the pressure (and the package) once you're over the **Install** window.

You'll see a progress indicator when your package is being installed.

6.4 Basic Configuration

You can customize Gnome-RPM through the **Preferences** dialog, which you can access from **Operations** => **Preferences** on the menu. To make selections in **Preferences** select the check boxes next to the options.

Under **Behavior**, you'll find a number of options for configuring the way Gnome-RPM installs, uninstalls and upgrades packages. The **Behavior** tab is split into five sections: Install, Upgrade, Other, Database and Architecture. (See Figure 6–4, *The Behavior Window* for an example.) Note that by default these boxes are not checked.

Figure 6–4   The Behavior Window
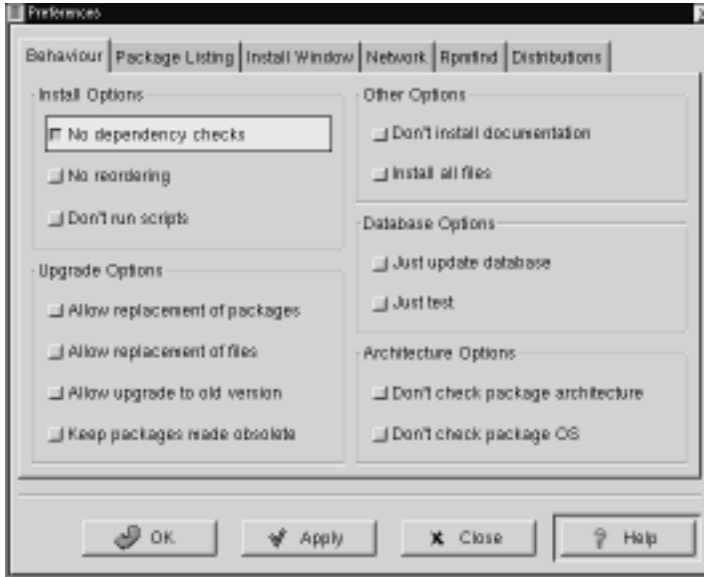
Under **Install Options**, you have the following choices:

• **No dependency checks** - When selected, this will install or upgrade a package without checking for other types of files on which the program may be dependent in order to work.

Please Note

Unless you know what you're doing, you should probably avoid using this option, since packages can be dependent on a vareity of programs or libries to properly function.

- **No reordering** - This option is useful if RPM is unable to change the installation order of some packages to satisfy dependencies.

- **Don't run scripts** - Pre- and post-install scripts are sequences of commands that are sometimes included in packages to assist with installation. This check box is similar to the --noscripts option when installing from the shell prompt.

Under **Upgrade Options**, you can select the following:

- **Allow replacement of packages** - Replaces a package with a new copy of itself. This option can be useful if an already-installed package has become damaged or may require other repair to function correctly.

- **Allow replacement of files** - Allows the replacement of files which are owned by another package. This option can be useful when there are two packages with the same file name but different contents.

- **Allow upgrade to old version** - This option allows you to "upgrade" to an earlier package -- useful if the latest version of an installed package doesn't correctly function for your system.

- **Keep packages made obsolete** - Prevents packages listed in an Obsoletes header from being removed.

In **Other Options**, you can select:

- **Don't install documentation** - Selecting this option can save on disk space by excluding documentation related to the package. Unless you know what you're doing, however, you're safest having the documentation which is included with an application, in case you require help with its features.

- **Install all files** - Installs all files in the package.

The choices available in **Database Options** and **Architecture Options** allow you to decide, among other things, whether you want to perform a "test" installation (which will

check for file conflicts without actually performing an install), or whether you want to exclude packages for other operating systems or system architectures.

In the **Package Listing** tab, you'll find a choice of displays for your packages: either as **icons** or as a **list**, which will provide more package information than the icon display.

In **Install Window**, you can specify the path through which Gnome-RPM can find new RPMs on your system (see Figure 6–5, *The Install Window*). The default path to serach for RPMs is

```
/mnt/cdrom/RedHat/RPMS
```

Note, however, that there is a selection of paths beneath the default choice, as well as a **Browse...** button next to the default path. The window listing the **RPM Directories** contains a listing of the most likely locations in which RPM files would be found on your system. By selecting the **Browse...** button, you can also specify an entirely different path. Once a new path is selected, click on the **Apply** button to save your default path for future sessions.

Figure 6–5   The Install Window
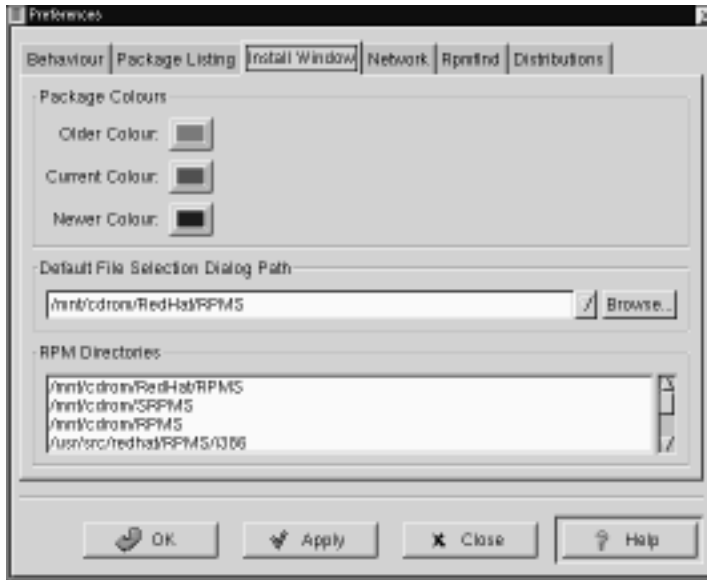
Optionally, if the path for your RPMs doesn't match the default path in your prefer-
ences, you'll be presented with a browser window, which will allow you to select the
correct path for your new RPMs.

In the **Install Window**, under **Package Colors**, you have the option of selecting the color
coding of your packages. The default color codes are set to gray for older packages,
green for current packages and blue for newer packages. Such color coding can be a
handy means of identifying both installed and uninstalled packages, as well as pack-
ages for which the Web find feature has located updates.

In addition to using the file dialog to add files to your install list, you can "drag" files
from the GNOME File Manager to add them to the list.

In **Rpmfind**, you have settings and options which correspond to the Web find feature.
The **Metadata server** selection sets the server to be used for searches. The **Download dir:**
entry allows you to specify where you want the files to be stored.

Figure 6–6   Distribution Settings in Preferences
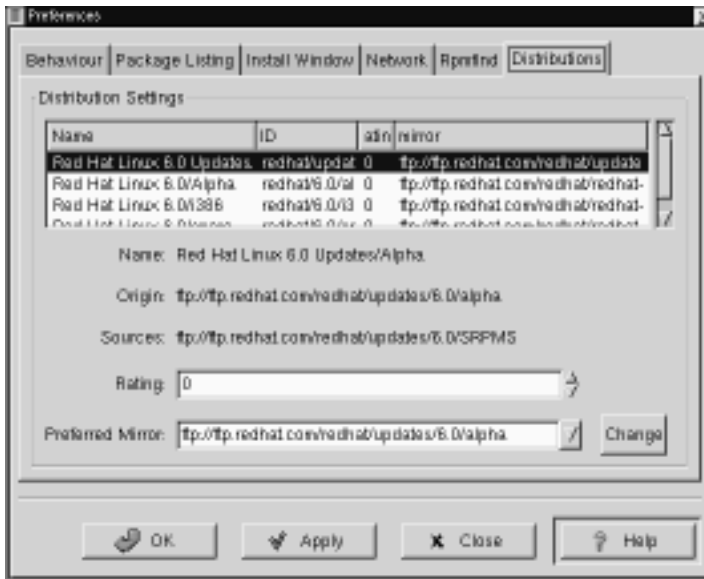
In **Distribution Settings**, you can set the options for choosing the most appropriate package out of the selections Rpmfind returns, as well as which mirror you would like to use. The higher the rating you indicate for your selection (as shown in Figure 6–6, *Distribution Settings in Preferences*) the higher the priority it will receive; the lower rating, such as "-1," will specify that packages not be recommended.

6.5 Package Manipulation

6.5.1 Querying Packages

The easiest way to query packages is to use the **Query** option from the menu at the top. If you want to query more than one package, make all your selections then press the **Query** button on the menu.

You'll be presented with a window like the one shown in Figure 6–7, *Query Window*. The more packages you've queried, the more "tabs" you'll find within the **Package Info** box, each tab representing a query response for the related package.

Figure 6–7   Query Window
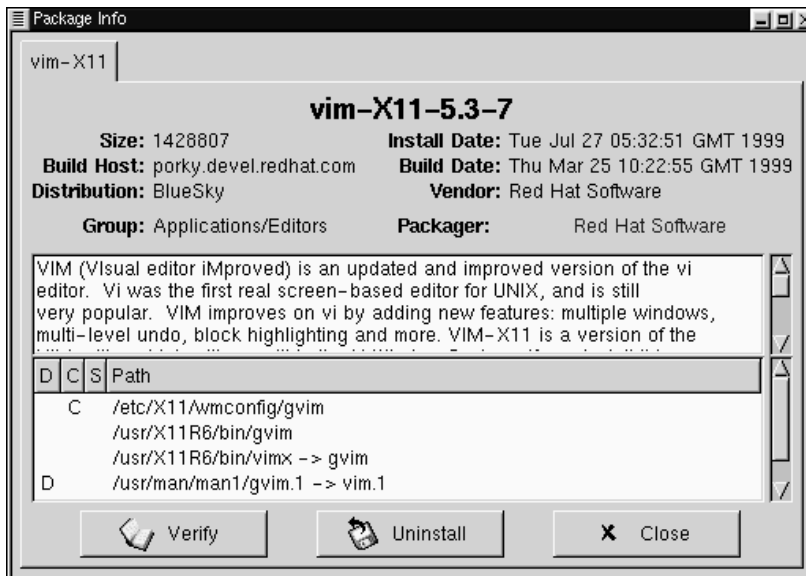
The name of the package is centered at the top of the box. Below, the box is divided into two columns of listed information; below this information, you'll see a display area showing package files.

In the left column of the information list, you'll find the size of the file, the machine on which the file is found, the name of the package distribution and the group to which its function belongs.

In the right column, you'll find the date of the package's installation on your machine, the date the package was built, the name of the vendor and the name of the group who

packaged the software. If the package has not been installed on your machine, that space will simply read, "not installed."

Below the description is a list of the files contained in the package. If a D appears in its related column to the left of the path, that file is a documentation file and would be a good thing to read for help on using the application. If a C appears in its respective column, the file is a configuration file. Under the S column, you can view the "state" of the package; here, you'll receive information if any files are reported as "missing" from the package (which probably means there's a problem with the package).

If you're querying a package that's already installed, you'll also find two additional buttons beneath at the bottom of this window: **Verify** and **Uninstall**. If you're performing a query on a package that hasn't been installed yet, the buttons on the bottom will be labeled **Install**, **Upgrade** and **Check Sig**.

To close the query window without performing any action, left-click on the **X** at the top right of the window bar.

6.5.2 Verifying Packages

Verifying a package checks all of the files in the package to ensure they match the ones present on your system. This check can be used when you suspect that one of the program's files has become corrupted for some reason.

6.5.3 Uninstalling Packages

Uninstalling a package removes the application and associated files from your machine. When a package is uninstalled, any files it uses that are not needed by other packages on your system are also removed. Changed configuration files are copied to *<filename>.*rpmsave so you can reuse them later.

Please Note

Remember that you must be root to uninstall packages.

If uninstalling a package would break "dependencies," a dialog will appear, asking you to confirm the deletion. This will occur if you haven't selected the **"No dependency checks"** box from the **Preferences** menu (as shown in Figure 6–8, *The Behavior Tab in Preferences*).

Figure 6–8 The Behavior Tab in Preferences
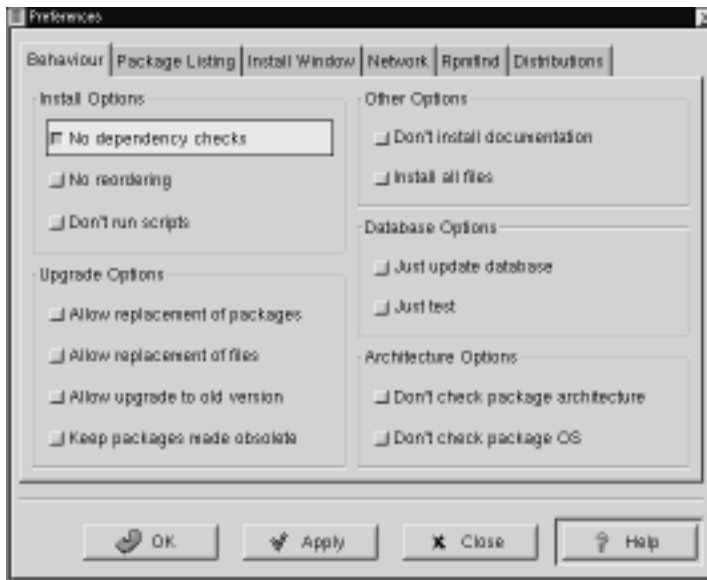
There are a variety of methods through which you can remove a selected package: from the menu, under **Packages**; from the toolbar and from the **Query** function.

Figure 6–9   Uninstall Window

Upgrading Packages

When a new version of a package has been released, it's easy to install it on your system; simply upgrade what you already have to the version that you want to install. The procedure is similar to installation of packages with the exception that you should select **Upgrade** rather than **Install**.

During the upgrade, you'll see a progress indicator similar to that for package installations. When the upgrade is finished, any old versions of the packages will be removed, unless you specify otherwise (refer to Section 6.4, *Basic Configuration* for more information).

It is much better to use the upgrade option than to uninstall the old versions of a package and then install the new one. (In fact, the RPM upgrade process will handle uninstallation and replacement of files. Refer to Section 6.6.2, *Uninstalling* to learn more about this process.)

6.6 RPM from the Shell Prompt

You can also use RPM technology by typing a vareity of commands at the shell prompt. At your disposal, of course, are all the functions available through Gnome-RPM, with some additions, such as the ability to **freshen** packages, which is similar to upgrading (refer to Section 6.6.5, *Freshening* for more information).

For more in-depth information about RPM, refer to the related chapters in the *Official Red Hat Linux Reference Guide*. You can also read the related man page by typing man rpm or the help file by typng rpm --help.

6.6.1 Installing

RPM packages typically have file names like foo-1.0-1.i386.rpm, which includes the package name (foo), version (1.0), release (1), and architecture (i386). Installing a package is as simple as typing:

```
# rpm -ivh foo-1.0-1.i386.rpm
foo                           ####################################
#
```

RPM prints out the name of the package (which is not necessarily the same as the file name), and then prints a succession of hash marks as the package is installed as a progress meter.

Actually, you can also install a package with the upgrade option. For more information, see Section 6.6.3, *Upgrading*.

While installing, you could run into errors, such as statements that the package is already installed or that there are **dependency problems** which hinder the installation of the package.

For information about the types of errors you may encounter, read on.

Package Already Installed

If the package is already installed, you will see:

```
# rpm -ivh foo-1.0-1.i386.rpm
foo                     package foo-1.0-1 is already installed
error: foo-1.0-1.i386.rpm cannot be installed
#
```

If you want to install the package anyway, you can use --replacepkgs on the command line, which tells RPM to ignore the error:

```
# rpm -ivh --replacepkgs foo-1.0-1.i386.rpm
foo                           ####################################
#
```

Conflicting Files

If you attempt to install a package that contains a file which has already been installed by another package, you'll see:

```
# rpm -ivh foo-1.0-1.i386.rpm
foo             /usr/bin/foo conflicts with file from bar-1.0-1
error: foo-1.0-1.i386.rpm cannot be installed
#
```

To cause RPM to ignore that error, use `--replacefiles` on the command line:

```
# rpm -ivh --replacefiles foo-1.0-1.i386.rpm
foo                             ###################################
#
```

Unresolved Dependency

RPM packages can "depend" on other packages, which means that they require other packages to be installed in order to run properly. If you try to install a package for which there is such an unresolved dependency, you'll see:

```
# rpm -ivh bar-1.0-1.i386.rpm
failed dependencies:
        foo is needed by bar-1.0-1
#
```

To correct the error, you should install the requested package. If you want to force the installation anyway (a bad idea since the package probably will not run correctly), use `--nodeps` on the command line.

6.6.2 Uninstalling

Uninstalling a package is also simple::

```
# rpm -e foo
#
```

Notice that when you uninstall a package, you should use the package *name* (`foo`) not the name of the original package ("`foo-1.0-1.i386.rpm`").

You can encounter a dependency error when uninstalling a package if some other installed package depends on the one you are trying to remove. For example:

```
# rpm -e foo
removing these packages would break dependencies:
        foo is needed by bar-1.0-1
#
```

To cause RPM to ignore that error and uninstall the package anyway (also a bad idea since the package that depends on it will probably fail to work properly), use `--nodeps` on the command line.

6.6.3 Upgrading

Upgrading a package is similar to installing.

```
# rpm -Uvh foo-2.0-1.i386.rpm
foo                                ###################################
#
```

When you upgrade, RPM automatically uninstalls any old versions of the `foo` package.

You can use `upgrade` to install packages, as well, since it works fine even when there are no previous versions of the package installed.

Since RPM performs intelligent upgrading of packages with configuration files, you may see a message like:

6.6.4 Errors on Installing Older Packages

```
# rpm -Uvh foo-1.0-1.i386.rpm
foo    package foo-2.0-1 (which is newer) is already installed
error: foo-1.0-1.i386.rpm cannot be installed
#
```

To cause RPM to "upgrade" anyway, use `--oldpackage` on the command line:

```
# rpm -Uvh --oldpackage foo-1.0-1.i386.rpm
```

```
foo                                   #####################################
#
```

6.6.5 Freshening

Freshening a package is similar to upgrading:

```
# rpm -Fvh foo-1.2-1.i386.rpm
foo                                   #####################################
#
```

RPM's freshen option checks the versions of the packages specified on the command line against the versions of packages that have already been installed on your system. When a newer version of an already-installed package is processed by RPM's freshen option, it will be upgraded to the newer version. However, RPM's freshen option will not install a package if no previously-installed package of the same name exists. This differs from RPM's upgrade option, as an upgrade *will* install packages, whether or not an older version of the package was already installed.

6.6.6 Querying

Querying the database of installed packages is accomplished with rpm -q. A simple use is rpm -q foo which will print the package name, version, and release number of the installed package foo:

```
# rpm -q foo
foo-2.0-1
#
```

Instead of specifying the package name, you can use the following options with -q to specify what package(s) you want to query. These are called *Package Specification Options*.

- -a queries all currently installed packages.
- -f *<file>* will query the package owning *<file>*.
- -p *<packagefile>* queries the package *<packagefile>*.

There are a number of ways to specify what information to display about queried packages. The following options are used to select the type of information for which you are searching. These are called *Information Selection Options*.

- -i displays package information such as name, description, release, size, build date, install date, vendor, and other miscellaneous information.
- -l displays the list of files that the package "owns".
- -s displays the state of all the files in the package.
- -d displays a list of files marked as documentation (man pages, info pages, README's, etc).
- -c displays a list of files marked as configuration files. These are the files you change after installation to adapt the package to your system (sendmail.cf, passwd, inittab, etc).

For those options that display file lists, you can add -v to your command line to get the lists in a familiar ls -l format.

6.6.7 Verifying

Verifying a package compares information about files installed from a package with the same information from the original package. Among other things, verifying compares the size, MD5 sum, permissions, type, owner and group of each file.

The command rpm -V verifies a package. You can use any of the *Package Selection Options* listed for querying to specify the packages you wish to verify. A simple use is rpm -V foo which verifies that all the files in the foo package are as they were when they were originally installed. For example:

- To verify a package containing particular file:

      rpm -Vf /bin/vi

- To verify ALL installed packages:

    ```
    rpm -Va
    ```

- To verify an installed package against an RPM package file:

    ```
    rpm -Vp foo-1.0-1.i386.rpm
    ```

    This can be useful if you suspect that your RPM databases are corrupt.

If everything verified properly there will be no output. If there are any discrepancies they will be displayed. The format of the output is a string of 8 characters, a possible "c" denoting a configuration file, and then the file name. Each of the 8 characters denotes the result of a comparison of one attribute of the file to the value of that attribute recorded in the RPM database. A single "." (period) means the test passed. The following characters denote failure of certain tests:

- 5 -- MD5 checksum
- S -- File size
- L -- Symbolic link
- T -- File modification time
- D -- Device
- U -- User
- G -- Group
- M -- Mode (includes permissions and file type)

If you see any output, use your best judgment to determine if you should remove or reinstall the package, or otherwise fix the problem.

7 Q & A: Quick Answers to Common Questions

7.1 Your First Login

7.1.1 Q: Logging In the First Time

I just installed Red Hat Linux, but now I see this screen that says "login." What do I do now?

7.1.2 A: Logging In as Root or User

When you installed Red Hat Linux, you were asked to create a **root password** and an account called **root** was created for you. The root account is all-knowing, all-powerful (at least, as far as your system is concerned).

You were also given the opportunity to create one or more **user accounts** at that time, for which you could specify a login name and a password.

When you see `login:`, then, type either the name and password of a user account you've already created during the installation, or log in as root by typing **root**.

In the `password:` field, type in the root password you were asked to create during the installation of Red Hat Linux, or the password for a user account.

What to do next? Well, if you're root -- the system administrator -- one of your first tasks should be to create a user account, if you haven't already done so. (Every good system administrator works within their user account, rather than rely on the root account for most mundane tasks.) Refer to Chapter 2, *Welcome to Linux* for instructions on how to create your user account.

7.2 Using a Diskette

7.2.1 Q: Working with Diskettes

How do I use a diskette with Red Hat Linux?

7.2.2 A. Using Mtools

If you've got an MS-DOS formatted diskette, you can access the files by using the mtools utility.

Mtools offers a wide range of options for working with diskettes, including copying, moving, deleting and formatting. To read more about mtools, open an Xterm window and type `man mtools` at the prompt.

For example, to copy a file from an MS-DOS formatted diskette (such as one from a Windows95 system), use the following syntax at the prompt in an Xterm window:

```
mcopy a:thisfile.txt thisfile.txt
```

The file will then be copied from the diskette drive (drive A:) to the directory you're currently in when you issue the `mcopy` command. If you're in your /home directory, then, you'll find `thisfile.txt` there.

If you want to view the contents of an MS-DOS formatted diskette, type `mdir` at the prompt. By default, you'll be shown the contents of the diskette in the A: drive.

To change to a subdirectory on the diskette, type

```
mcd a:subdir
```

(The above command assumes `subdir` is the name of the subdirectory you want to access.)

If you have a diskette which you want to format and use with your Red Hat Linux system, use the **ext2 filesystem**. For more on this, read on…

7.2.3 A: The ext2 Filesystem

To use a diskette specifically with Red Hat Linux, you'll need to create a **Second Extended (ext2) filesystem** on the disk. The ext2 filesystem is the filesystem used by Red Hat Linux, and is the most commonly used Linux filesystem type.

Once you've created an ext2 filesystem on the diskette, you can manipulate its contents in the same ways that you manipulate directories and files on your hard drive.

7.2.4 Creating an ext2 Filesystem on a Diskette

The mke2fs command is used to create a Linux ext2 filesystem on a device, which may be a disk partition or a diskette. Put your formatted diskette into the drive and issue the following command at a shell in an Xterm window:

```
$ /sbin/mke2fs /dev/fd0
```

On Linux systems, /dev/fd0 refers to the first diskette drive, usually your A: drive.

The mke2fs utility has a number of options. The -c option makes the mke2fs command check the device for bad blocks before creating the filesystem. The other options are covered in the mke2fs man page.

Once you've created an ext2 filesystem on the diskette, it is ready to be used with your Red Hat Linux system.

7.3 Switching Environments

7.3.1 Q: Options for Switching Environments

I know I have both KDE and GNOME on my Red Hat Linux system, but how do I change from one environment to the other?

7.3.2 A: Using `Switchdesk`

You can use a utility called **Switchdesk**, which allows you to change to and from GNOME, KDE and other environments.

Figure 7–1   The Switchdesk Utility from the Shell Prompt
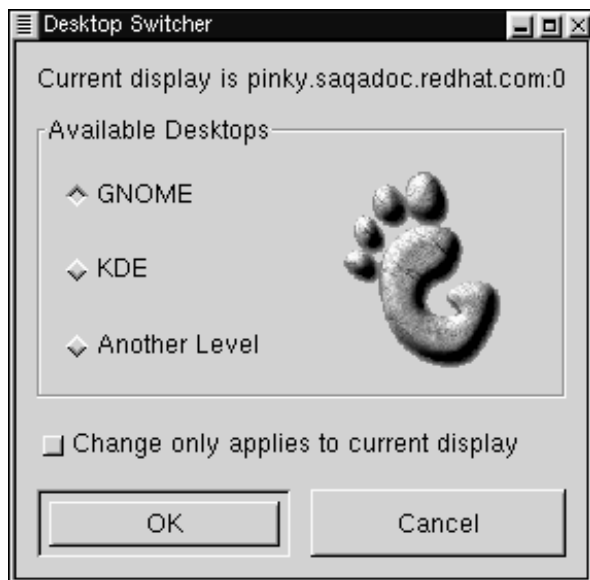
You can use Switchdesk from either the shell prompt or from the login screen.

From the shell prompt: Open an Xterm window and type

```
switchdesk
```

A window will appear (see Figure 7–1, *The Switchdesk Utility from the Shell Prompt*) showing you the choices of environments you can select, under **Available Desktops**. Once you select your new environment, click on **OK**.

To see your changes take effect, you'll have to log out, then log back in.

From the login screen: You'll select other alternative environments under **Options** => **Sessions**. Again, you can select a variety of environments.

Now, just type in your name and password, and log in as usual.

7.4 Shutting Down

7.4.1 Q: "Right way" to Shut Down

What's the "right way" to shut down or reboot my system? Are there any shortcuts?

7.4.2 A: Putting Away Files

Any method you choose which allows your Red Hat Linux system to put away all its data files and stop running processes is the "right way" to shut down. *Never* simply turn off your machine.

Here are two ways you can cleanly shut down your system:

- From your X session: Go to the **Panel's Main Menu Button** => **Log out** then choose **Logout**, **Halt** or **Reboot**. You can also choose to save your current setup, which means that the programs which are running when you log out will return when you log back in.

 If you select either **Halt** or **Reboot**, you'll be required to enter your login password after you press the **OK** button to verify your selection.

- From the shell prompt: In an Xterm window, type shutdown -r now or shutdown -h now. The -r means "reboot" and the -h means "halt." You will be required to enter your password before these commands are executed.

- From the login screen: Go to **Option** => **System**. You can choose whether to halt or reboot your system, without a required password.

A Shortcut to Rebooting

One of the fastest ways to reboot your system is to perform the "three-fingered salute" -- Pressing the [Ctrl]-[Alt]-[Del] keys all at the same time -- from a **virtual console**.

A virtual console is a character cell interface, unlike the graphical interface of the X Window System. There are several virtual consoles available to you, but the easiest to remember is probably the first one. You can reach it by pressing [Ctrl]-[Alt]-[F1].

Getting Back

To return to your X session from a virtual console, just press
the [Ctrl]-[Alt]-[F7] keys.

To reboot from a virtual console, first move to a virtual console by pressing [Ctrl]-[Alt]-
[F1], then press [Ctrl]-[Alt]-[Del] to reboot.

7.5 Error Message While Installing an RPM

7.5.1 Q: Errors on Software Installation

How do I install an RPM from a CD or the Internet? I keep getting an error message
when I use Gnome-RPM.

7.5.2 A: Having "Permission"

If you're getting the error message similar to "failed to open /var/lib/rpm/pack-
ages.rpm" (see Figure 7–2, *An RPM Install Error*) it's likely due to the your
permissions.

When you install software, you're often required to make system-wide changes --
changes which only root can enact. In your user account, permission to make such
changes will be unavailable to you by default.

Figure 7–2   An RPM Install Error
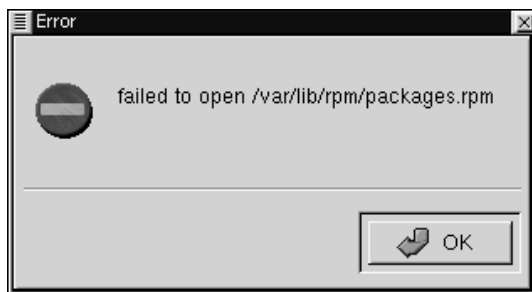
Here's a way to install RPMs in your home directory -- so you can work with them in your user account.

1. Open an Xterm window and, at the shell prompt, type su, then your ***root-password***.

2. At the prompt -- which should now be a hash mark (#) -- type gnorpm to start Gnome-RPM.

3. Install or upgrade your package.

4. Exit from Gnome-RPM then, to exit from **superuser** mode, at the shell prompt, type exit.

For more information about RPM and Gnome-RPM, refer to the chapters in this guide or, for greater depth, refer to the *Official Red Hat Linux Reference Guide*.

7.6 Starting Applications

7.6.1 Q: How Can I Start an Application I Downloaded?

I installed an application I downloaded, and everything seemed to go fine, but I still get "command not found" when I type its name. I think I have the right name, so why won't it start?

7.6.2 A: Starting Applications.

If you're trying to start an application from the shell prompt and it's proving fruitless, try preceding the name of the application's executable with a ./.

Let's say you've downloaded a setiathome client and want to try it out. You follow the directions for installing the software. Now you change to the directory in which you know the exeutable can be found (as shown below).

```
cd setiathome
```

To start the application now, precede the executable with a "./," as shown below:

```
./setiathome
```

Briefly, the reason you've got to use the `./` in order to start the application is because the executable wasn't placed in a directory where your shell knew it could be found (such as `/usr/bin`).

In such instances, you often have to go into the directory which holds the executable and start the application from there. That means you'll have to tell your shell where it can find the executable -- stating `./` tells bash the executable can be found in the "current working directory."

You can customize your settings so that you won't be required to use the `./`, however. See the following for more info on how to accomplish this.

Editing Your PATH

If you want to periodically start programs without having to enter a `./` before the executable, you'll have to do a little tweaking.

Basically, you'll have to add the "current working directory" (.) to your the list of directories in your **PATH** environment variable, letting the shell know that it can start applications in whichever directory you're currently working.

CAUTION

These instructions are intended *only* for user accounts. Avoid modifying files such as `.bash_profile` for root, because of the potential security risks.

Using a text editor, such as pico, in an Xterm window, open the file called `.bash_profile` by typing

```
pico .bash_profile
```

You'll see a **PATH** statement, similar to the one shown below.

```
PATH=$PATH:$HOME/bin:/usr/lib/
```

Now, to the end of this statement, add **:./**, as shown below

```
PATH=$PATH:$HOME/bin:/usr/lib/:./
```

Now, type [Ctrl]-[X]; you'll be asked whether you want to save "the modified buffer"; type [Y] for "yes." Next, you'll see the name the file will be saved as; press the [Enter] key.

You're done. Now, you won't be required to start some applications with the ./ preceding the name of the executable.

7.7 Accessing a Windows Partition

7.7.1 Q: How Do I Access My Windows Partition?

I've got a dual-boot system with Red Hat Linux and Windows 98. Is there a way I can gain access to my Windows partition while I'm running Linux?

7.7.2 A: Two Ways to Access Windows Partition.

You can gain access to another partition on your system -- for example, a Windows partition -- by typing a line similar to the following:

```
mount -t vfat /dev/hda1 /mnt/vfat
```

The foregoing assumes that your Windows partition is on your first hard drive, in the first partition (/dev/hda1). Also, before typing the above command, as root, you should create a **mount point** called /mnt/vfat (Using the command mkdir /mnt/vfat) to hold the data.

Another method of mounting a Windows partition is by entering the correct fields in the /etc/fstab file. One of the easiest ways to enter these fields is through Linuxconf.

Open an Xterm window. If you're in your user account, su to root by typing

```
su
Password:yourrootpassword
```

Now, create a mount point for your Windows partition, by typing

```
mkdir /mnt/vfat
```

Next, while you're still root, start Linuxconf by typing `linuxconf` at the prompt.

Scroll down in the "tree view," or left panel of Linuxconf, to the entry marked **File Systems**, and click on the **+** to expand the tree.

Now, click on the entry **Access local drive**. In the right panel, you'll see a list of the currently mounted filesystems. (Your new mount point isn't there because you haven't added it yet. That's coming up…)

Click on the **Add** button. A tab, marked **Volume specification** will appear. In the first tabbed entry, called **Base**, you'll be presented with a series of boxes to fill out, either by typing or by selecting from a drop-down list. Here's what the boxes represent:

> Partition: The physical location of your Windows partition (for example, `/dev/hda1` for the first partition on the first hard drive);

> Type: The filesystem type. A Windows 98 partition, for example, would be `vfat`;

> Mount point: The name of the mount point you'd chosen earlier (for example `/mnt/vfat`).

There are a few other options, as well, from various tabbed windows; these concern setting permissions for allowing users to access the partition, choosing whether to *prevent* the partition from being mounted at boot time, and other choices. Some options to pay particular attention to include:

> "Dos options": Both "default user id" and "default group id" should be set to your user account if you want to be able to access while you're logged in as user (rather than having to `su` to root, which is the default mode); "translation mode" should be set to "AUTO" for most purposes, rather than specify either "text" or "binary"; "default permission" specifies whether those with access to the partition can read, write to and/or execute files and directories on the partition. You might select a setting of **755**, which allows read, write and execute for the user, and read and execute permissions to the group and others.

> "Options": Unless you specify that you do not want the partition to be mounted when you boot your system, it isn't necessary to select the "User mountable" option in this tab.

Make your selections -- if you become stuck, choose the **Help** button for additional pointers.

When you're finished, select the **Mount** button to check your new partition. You should see verification that the partition has been successfully mounted.

Now, click on the **Accept** button. Your new filesystem will be added to /etc/fstab.

To acces the partition now, type cd /mnt/vfat, and to navigate through Windows 98's "long filename" directories, surround the directory in quotation marks, as in ls "Program Files".

7.8 Find Some Commands Fast

7.8.1 Q: Locating Previously Used Commands

I was looking at a man page yesterday, but I can't remember the name of the command I was reading about, and I didn't write it down. How do I get the man page back?

7.8.2 A: Seaching .bash_history.

Chances are good the command you used is stored away in a file called .bash_history. By default, this file stores away the last 500 commands you typed at the shell prompt.

You can glimpse the history of your commands by typing history at the shell prompt -- but the results will speed by quickly.

Another way you can view .bash_history is with a **pager** such as less, by typing less .bash_history at the shell prompt. To move forward a screen, press the [Spacebar]; to move back a screen, press the [B] key, and to quit, press [Q].

But paging through to find a command can be tedious. A useful alternative is to search the file for keywords using grep, a powerful search utility.

If you can recall part of any command you'd typed at the shell prompt, and if that command is still in .bash_history, you can find it fast by **piping** the results of the history command through the utility called grep.

Let's say you'd been reading the man page the day before, but can't recall its name. To search for the command, then, type:

```
history | grep man
```

In no time, you'll see all of the commands you'd typed which have the word *man* in them -- both the relevant and irrelevant.

There are plenty of ways to make use of command history. For other tips and tricks, see Section 7.10, *Tips on Using Command History*

7.9 Keep `ls` Output from Scrolling

7.9.1 Q: When `ls` Output Scrolls Off the Screen

Whenever I type `ls` I can barely see the output of the directory because it scrolls by too quickly. How can I actually read the output?

7.9.2 A: "Piping" the Output of `ls`

You can prevent the output of `ls` from scrolling by too quickly by **piping** the output to a pager, such as `less` or `more`. The results are akin to the DOS command `dir` *somedirectory* /p, because you'll see the output one screen, or "page" at at time.

To read the output of `/etc` with `less`, for example, at the shell prompt, type:

```
ls -al /etc | less
```

To move forward a screen, press [Space]; to move back a screen, press the [B] key; to quit, press [Q].

You can achieve the same results with `more`, another pager.

Printing Output of `ls`

You can also print the directory by piping the output to a printer in the same way that you piped the output to your screen. Assuming you've configured a printer, to pipe the output type:

```
ls -al /etc | lpr
```

7.10 Tips on Using Command History

7.10.1 Q: Tips and Tricks for `history`

What are some other ways I can use command history?

7.10.2 A: Using More Command History

If you type `history`, you'll see a numbered list flash by, showing you the previous 500 commands you'd used.

You probably don't need to see all of the last 500 commands, so the command `history 20` might be more useful -- where `20` is the previous 20 commands you'd typed. (You can use any number to reveal a specific number of previous commands.)

Other Shortcuts

Here are other command history shortcuts which may come in handy for you:

- "Bang, bang": Typing `!!` (called "bang bang") executes the last command in the history.

- "Bang *number*": Typing `!number` (as in `!302`) will execute the command which is numbered 302 in the history file.

- "Bang *string*": Typing `!string` (as in `!rpm`) will execute a command with that matching string from the history file.

- [Up arrow] and [down-arrow]: At the Bash shell prompt, you can simply press the [Up arrow] to move back through previous commands in your history list (the [Down arrow] will move you forward through the commands) until you find the command you want. Press [Enter] to execute the command, just as if you had typed it on the command line.

7.11 Unable to Log In

7.11.1 Q: Using `linux single`

Help! I can't remember my root password. How do I log in now?

7.11.2 A: Use Single-User Mode

You can log in with **single-user mode**, in which you can create a new root password.

When you see the `boot:` prompt, type **linux single** to enter single-user mode. Some of the filesystems will be mounted, and you will find a `bash#` prompt when

you've entered single-user mode (note that this prompt will look somewhat different than the prompt you're accustomed to).

Now, you can change root's password by typing

```
bash# passwd root
```

You'll be asked to re-type the password for verification. Once you're finished, the password will be changed and you can reboot by typing shutdown -r now at the prompt; then you can log in to root as before.

7.12 Starting "Services"

7.12.1 Q: Starting Services at Boot Time

I'd like to see which services are automatically started on my system, and whether I can add any. How do I do that?

7.12.2 A: Use `ntsysv`

By using ntsysv, you can both see which services are automatically started, and add to the list -- although you're probably better off leaving the default values, unless you know what you're doing.

To start the utility, type ntsysv at the shell prompt. You should be root when you use the utility, however, because you're making modifications to your system.

Among the services you'll find are sendmail, the powerful e-mail utility, httpd, the Apache World Wide Web server, nfsfs, which mounts and unmounts NFS file systems.

If you have questions about the services, you can press the [F1] key to obtain information about the service.

7.13 Using All Available RAM

7.13.1 Q: Using All the Memory

My computer has 128 MB of RAM, but I'm not sure all of it is being used. How can I find out whether all the RAM is being used? Is there a way to get Linux to recognize all my memory, if it isn't already?

7.13.2 A: Instruct the Kernel to Use All Your RAM

Usually, the Linux kernel will recognize all of the RAM on a system; however, there may be circumstances in which not all of the memory is being recognized.

There are a couple ways to make certain that all of your machine's memory is being used by the Linux kernel. First, you can test to see if it's being used and, if it isn't, you can either enter the option at the `boot:` prompt or you can edit a file called `/etc/lilo.conf` to make sure that the memory is used each time you boot into your system. (The `/etc/lilo.conf` is the configuration file for LILO, the LInux LOader.)

First, open an Xterm window and, at the shell prompt, type `cat /proc/meminfo` to find out the current statistics about memory usage on your system. Your output will probably look similar to the following:

```
        total:    used:    free:  shared: buffers:  cached:
Mem:  64655360 63033344  1622016 51159040  1552384 33816576
Swap: 73990144   700416 73289728
MemTotal:     63140 kB
MemFree:       1584 kB
MemShared:    49960 kB
Buffers:       1516 kB
Cached:       33024 kB
SwapTotal:    72256 kB
SwapFree:     71572 kB
```

Among the categories you'll find will be the total amount of memory the Linux kernel "sees" (`Mem:`). If, for example, you have 128 MB of RAM, and you see the `Mem` entry reporting only 64 MB or so, then you can instruct the kernel to use all the RAM on your system.

You can pass the instruction to the kernel at the LILO boot: prompt by typing it in, such as

```
boot: linux mem=128M
```

Once you press [Enter], you'll boot with your new memory options. If you don't want to type the option each time you boot, you can enter the option in /etc/lilo.conf. You can use the editor called Pico.

But first, you'll have to su to root. In an Xterm, type

```
[you@localhost you]$ su
password: yourrootpassword
[root@localhost you]#
```

Now, to edit /etc/lilo.conf, in the same Xterm window, type

```
[root@localhost you]# pico /etc/lilo.conf
```

The file will probably look similar to the following:

```
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
image=/boot/vmlinuz-2.2.5-15
        label=linux
        root=/dev/hda1
        read-only
```

In this file, place the following as the first line: **append="mem=128M"**. With the addition, then, the file should look something like this:

```
append="mem=128M"
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
image=/boot/vmlinuz-2.2.5-15
        label=linux
```

```
        root=/dev/hda1
        read-only
```

You can use the [Ctrl]-[X] to quit the file.

You'll be asked if you want to save the file; select [Y] to save your changes, [N] if you don't want to save the changes, and [Ctrl]-[C] to cancel.

Now, type /sbin/lilo at the prompt to restart LILO. You can type exit to leave root, and type exit again to close the window.

When you reboot, your memory configuration will be passed to the kernel.

Part II   The GNOME User's Guide

8 The GNOME User's Guide

8.1 What is GNOME

GNOME is a user-friendly desktop environment that enables users to easily use and configure their computers. GNOME includes a panel (for starting applications and displaying status), a desktop (where data and applications can be placed), a set of standard desktop tools and applications, and a set of conventions that make it easy for applications to cooperate and be consistent with each other. Users of other operating systems or environments should feel right at home using the powerful graphics-driven environment GNOME provides.

GNOME is completely open source (free software), with freely available source code developed by hundreds of programmers around the world. If you would like to learn more about the GNOME project please visit the GNOME web site at http://www.gnome.org.

GNOME has a number of advantages for users. GNOME makes it easy to use and configure applications without using text-only interfaces.

GNOME is highly configurable, enabling you to set your desktop the way you want it to look and feel. GNOME's session manager remembers previous settings, so once you've set things the way you like they'll stay that way. GNOME supports many human languages, and you can add more without changing the software. GNOME even supports several Drag and Drop protocols for maximum interoperability with applications that aren't GNOME-compliant.

GNOME also has a number of advantages for developers, which indirectly also help users. Developers don't need to purchase an expensive software license to make their commercial application GNOME compliant. In fact, GNOME is vendor neutral - no component of the interface is controlled solely by one company or restricted from modification and redistribution. GNOME applications can be developed in a variety of computer languages, so you're not stuck with a single language. GNOME uses the Common Object Request Broker Architecture (CORBA) to allow software components to interoperate seamlessly, regardless of the computer language in which they

are implemented, or even what machine they are running on. Finally, GNOME runs on a number of Unix-like operating systems, including Linux.

GNOME is an acronym for the GNU Network Object Model Environment, so GNOME is a part of the larger GNU project. The GNU Project started in 1984 to develop a completely free Unix-like operating system. If you'd like to learn more about the GNU project you can read about it at http://www.gnu.org.

8.2 About This Guide

This users guide is designed to help you find your way around GNOME with ease. Both new and experienced computer users can benefit from this guide. If you're new to GNOME, or even computers, you'll gain an idea of how to use your desktop. If you're an advanced computer user trying out GNOME, you can work with expert tips, which will help you to become familiar with GNOME.

Although this was written originally in English, there are many translations of the guide available now or in the near future. If you would like to have this guide in another language you should check your operating system distribution or visit the GNOME Web Site to find out more information on translation.

9   GNOME Quick Start

9.1 GNOME Quick Start

Figure 2-1 shows an example of GNOME running. GNOME is very configurable, so
your screen may look quite different.

Figure 9–1   Sample GNOME Display.

The long bar at the bottom of *Figure 2-1* is a GNOME Panel, which contains a collec-
tion of useful panel applets and menus. Panel applets are tiny programs designed to
be placed in a panel, for example, the clock applet on the far right shows the current
time. The arrows on each side of the panel hide (and unhide) the panel.

The button in the panel containing a stylized foot is the Main Menu Button. Just
click on the Main Menu Button and you'll see a menu of pre-loaded applications and
actions, including a logout command.

The rest of the screen space is called the "Desktop." Just place on your desktop the items you use most often and you can double-click on an item (with the left mouse button) to use it:

- If the item is a program, that program will start.
- If it's data, the appropriate program will be start up with that data loaded.
- If it's a directory, the file manager application will start and show the contents of that directory. Your desktop will probably have a folder icon labelled "Home directory". Double-clicking on it will start a file manager at your home directory.

The file manager application lets you manipulate your files. The left side of its window shows directories, and the right side shows the selected directory's contents:

- To move the file or directory, just drag and drop it.
- To copy a file, hold down the [CTRL] key while dragging.
- To run a program or edit a data file, double-click it.
- To perform other operations on a file (such as rename or delete), select it using the right mouse button.
- To select more than one item at a time, click on the items after the first one while holding down the SHIFT key.

You can easily move or copy files between directories by starting two file manager applications, each one showing a different directory. If you want to put a file on your desktop, simply drag it from the file manager onto the desktop. In fact, dragging and dropping items onto other items generally "does the right thing" in GNOME, making it easy to get work done.

GNOME is very configurable; for example, you can have multiple panels (horizontal and vertical), choose what goes in them, and have them hide automatically. There are many panel applets you can include in your panel. You can also change how the screen looks; later portions of this document tell you how.

GNOME follows several UNIX conventions you should be aware of. The left mouse button is used to select and drag items. The right mouse button brings up a menu for the selected object (if a menu applies). Most UNIX mice have 3 buttons, and

the middle button is used to paste text (if in a text area) or to move things. If you only have two buttons, press the left and right buttons simultaneously to simulate the middle button. To copy text, use the left button to drag across the text you want to copy, move to the place you want the text to be, and press the middle button.

When an application window is displayed, there will be some buttons in its borders for controlling the window. These include buttons to minimize, maximize, and close the window. Their appearance can be configured and is controlled by a component called the "window manager."

Two examples of border styles are the Clean style (*Figure 2-2*) and the ICE style (*Figure 2-3*):

- In the Clean border style, the underscore means minimize, the square means maximize (use the whole screen), and the X button means close the window.

- In the ICE style, the X button will close the window. Clicking the arrow with the left mouse button minimizes the window, while clicking with the right mouse button shows a menu of other options.

Figure 9–2   Clean Border Style

Figure 9–3   ICE Border Style

If you are using a default installation of GNOME you may notice that minimizing a window actually causes that window to disappear from your desktop. To regain that window you may use the GNOME Pager, which is located on the Panel. The Pager will show you which tasks are running and where they are on your desktops. You will find the application you minimized in the task list on the right side of the GNOME Pager. Press the button for that application and it will return to your desktop.

Figure 9–4   The GNOME Pager
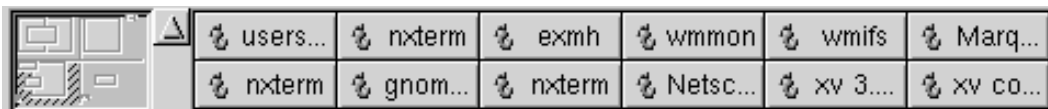
You may read more about the GNOME Pager in Section 19.6.4, *GNOME Pager*.

The following sections go into more detail, describing each component of the system.

10   Window Managers and GNOME

10.1 About Window Managers

The window manager is the piece of software that controls the windows in the X window environment. The placement, borders, and decorations of any window are managed by the window manager. This is very different from many other operating systems, and the way GNOME deals with window managers is different from other desktop environments.

As stated earlier in this guide, GNOME is not dependent on any one window manager. This means that major parts of your desktop environment will not change when you decide to switch window managers. GNOME works with the window manager to give you the easiest work environment you can have. GNOME does not worry about window placement but gets information from the window manager about their placement. The GNOME Pager will only work with a GNOME compliant window manager as will drag and drop on the desktop.

At the time of this version of the GNOME User's Guide the Enlightenment Window Manager is the only window manager that is 100% compliant. There are many other window managers that are partially compliant or are being modified to meet compliance.

Some of the window managers that have partial to full compliance at the time of this version of the GNOME User's Guide are:

*   Enlightenment - http://www.enlightenment.org

*   Icewm - http://www.kiss.uni-lj.si/~k4fr0235/icewm/

*   Window Maker - http://www.windowmaker.org

*   FVWM - http://www.fvwm.org/

*   AfterStep - http://www.afterstep.org/

There are a host of newer window managers being developed that will work with GNOME. You can find a list of these on the GNOME Software Map.

10.2 Changing Window Managers

At any time you may change the window manager you are using by utilizing the Window Manager Capplet in the GNOME Control Center. You may read more about this Capplet in Section 17.2.4, *Window Manager Capplet*

IMPORTANT

Keep in mind that the window manager you choose to use may not be compliant with GNOME and you may not benefit from some of the GNOME features if you use it.

11   The Enlightenment Window Manager

11.1 Introduction

GNOME is a desktop environment that is not dependent on any one window manager in order to control how applications and other features interact with each other. GNOME can work with a variety of window managers. The window manager is what controls the window borders, window decorations, and the functionality of those windows. Currently there are about three or four window managers that are considered GNOME compliant. Being GNOME compliant means much more than simply being able to run with GNOME; it also means being aware of GNOME and what GNOME offers. This includes session management, desktop settings, and interactivity with applications such as the GNOME Pager.

The Default Window Manager for this release of Red Hat Linux is a window manager called Enlightenment. Within Enlightenment there is a tool which will help you to configure the window manager to your specifications.

11.2 The Enlightenment Configuration Tool

There are two ways to launch the Enlightenment Configuration tool (E-conf). You can open the GNOME Control Center and select the **Window Manager Capplet** from the **Desktop** group. If you have Enlightenment as your current window manager you can run E-conf by pressing the **Run Configuration Tool for Enlightenment** button.

You can also launch E-conf by clicking your middle mouse button (or both right and left mouse buttons if you have emulation) on the desktop and selecting the **Enlightenment Configuration** item from the **pop-up** menu.

E-conf is divided into many sections that control various effects and functions within Enlightenment (E). You will find a list of these sections in the top left corner of the E-conf. The sections to choose from include: Basic Options, Desktops, Behavior, Audio, Special FX, Backgrounds, Themes, and Shortcuts.

11.2.1 Basic Options

The Basic Options section of E-conf lets you configure how E controls the movement and focus of windows on your desktop. Each method for controlling windows has a small image to show you what they look like, but perhaps the best way to get a feel is to select one, press the **Apply** button and move the window around. There are three areas to configure: Move Methods, Resize Methods, and Keyboard focus follows.

Figure 11–1 Enlightenment Basic Options
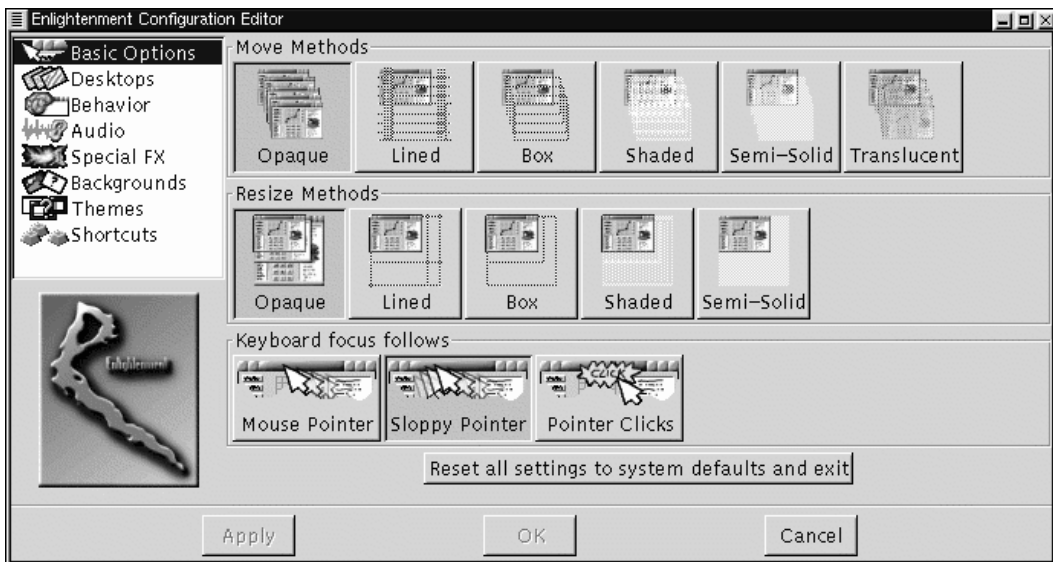
- **Move Methods** - This area controls the way windows look when you move them around your desktop with your mouse. You can choose from Opaque, Lined, Box, Shaded, Semi-Solid, and Translucent. Depending on your machine speed some of these move methods may work slower than others.

- **Resize Methods** - The Resize Methods option controls how the windows look when you resize them. You can choose among Opaque, Lined, Box, Shaded, and Semi-Solid. Again, depending on your machine speed some of these methods may work slower than others.

- **Keyboard focus follows** - This area controls which window gets the focus. Getting the focus means that the keyboard can control the window and if there is a distinction in you enlightenment theme, the window will be highlighted. You can choose from Mouse Pointer, Sloppy Focus, and Pointer Clicks.

  Mouse Pointer will focus the window whenever the mouse is over the window. **Sloppy Pointer** will do the same but will focus the window when the mouse is over it and keep the focus until the mouse is over another window. For example, you can put the mouse over a window and then over the desktop and the window will still have focus until it is over another window. **Pointer Clicks** means that you have to actually click on the window to give it focus.

At any time you can reset all of the Enlightenment configurations to the default settings by pressing the **Reset all settings to system defaults and exit** button.

11.2.2 Desktops

As you read in Section 13.3, *Desktop Areas*, GNOME supports the use of multiple desktop areas, and for the advanced user, multiple desktops. The control over these areas and desktops is actually handled by the window manager. The Enlightenment Configuration tool allows you to change the number of areas and desktops you have use of in your work environment.

Figure 11–2   Enlightenment Desktops
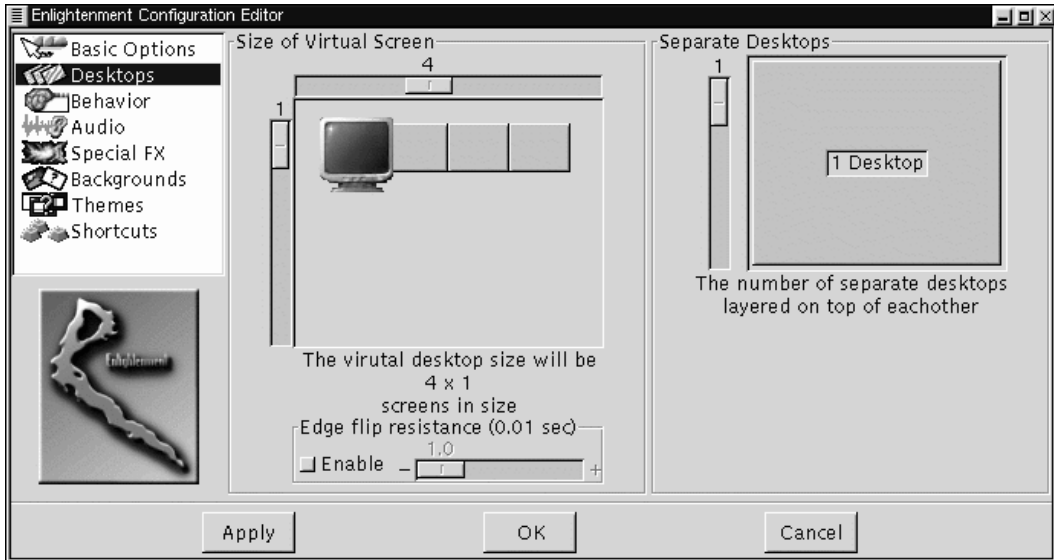
The Desktops Section of E-conf has two main tools to let you change the number of areas and desktops you use. The first tool is the **Size of Virtual Screen** which controls the number of desktop areas. Desktop areas are really just one big screen broken into many "virtual areas". The advantage to using virtual areas over multiple desktops is that drag and drop on the desktop will work across areas as will work being done in applications such as the GNOME File Manager. To use this tool slide the two sliders until you have the number of virtual areas, across and down, that you wish to use.

The second tool is **Separate Desktops**, which allows you to set the number of desktops you wish to use. Unlike desktop areas these are completely separate desktops. To use this tool move the slider up or down until you have the number of desktops you wish to use.

One good way to visualize your multiple areas/desktops situation is to make sure you have the GNOME Pager running. The GNOME Pager will show you how many areas and desktops you have and where your applications are within those spaces. You can

also use the GNOME Pager to navigate from area to area and desktop to desktop. You can read more about the GNOME Pager in Section 19.6.4, *GNOME Pager*.

Edge resistance factor for auto flip - If you are using desktop areas you can choose to have the areas switch when your mouse goes past the edge of the screen in the direction of another area. If you do have this feature on you can change the resistance the mouse has when it is moving past the edge by moving the slider to the number of seconds it takes to change to the next desktop.

11.2.3 Behavior

The Behavior Section of E-conf allows you to control various focus and miscellaneous behaviors in E.

Figure 11–3 Enlightenment Behavior
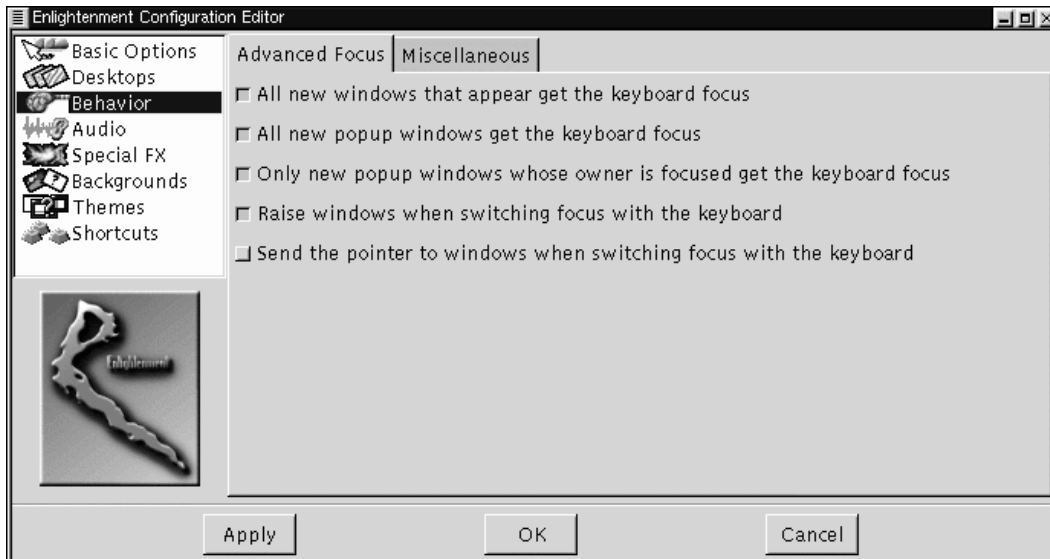
- **Advanced Focus** - This tab allows you to change properties related to focusing windows.

All new windows that appear get the keyboard focus - This ensures that when you start a new application you can start working on it when it appears without moving the mouse to get the focus of the window.

All new pop-up windows get the keyboard focus - This will give the focus to new pop-up windows only. If you have this selected and the "All new windows that appear get the keyboard focus" not selected, your pop-ups will get the focus no matter where they appear, but new applications started will not. An example of a pop-up window is an error message you might get in an application.

Only new pop-up windows whose owner is focused get the keyboard focus - This will give the focus to new pop-up windows only if the application that launched the window has the focus beforehand.

Raise windows when switching focus with the keyboard - If you use keystrokes to change the focus of your windows this will bring the focused window to the foreground.

Send pointer to windows when switching focus with the keyboard - If you use keystrokes to change the focus of your windows, this will send your mouse pointer to the focused window.

- **Miscellaneous** - This tab has miscellaneous Enlightenment behaviors.

Tooltips ON/OFF & timeout for tooltip pop-up - This controls whether you would like tooltips to display on various parts of an Enlightenment-controlled window. A tooltip is a small information window which opens when your mouse pointer hovers over an element in an application or component. The time slide bar next to the button lets you change how long it takes for the tooltips to appear once your mouse is over a component controlled by E. The speed is measured in seconds.

Automatic raising of windows after X seconds - This will focus a non-focussed window after the seconds you specify when your mouse is over it.

Transient pop-up windows appear together with leader - If your application brings up small pop-up windows as part of its normal function this will ensure that those pop-up windows stay with the main window. This can be helpful if your work flow is interrupted by the pop-up window showing up on the opposite side of your screen while you are working.

Switch to where pop-up window appears - This will send the mouse cursor and keyboard control to any pop-up window that appears on the desktop.

Display icons when windows are iconified - If you check this option and minimize your window this will display available icons for the application running in the window onto a tile that is created. If you are using the GNOME Pager you do not need to enable this option as it will allow you to access minimized applications.

Place windows manually - If you wish to place a new window yourself this will give the mouse the placement control of any new window.
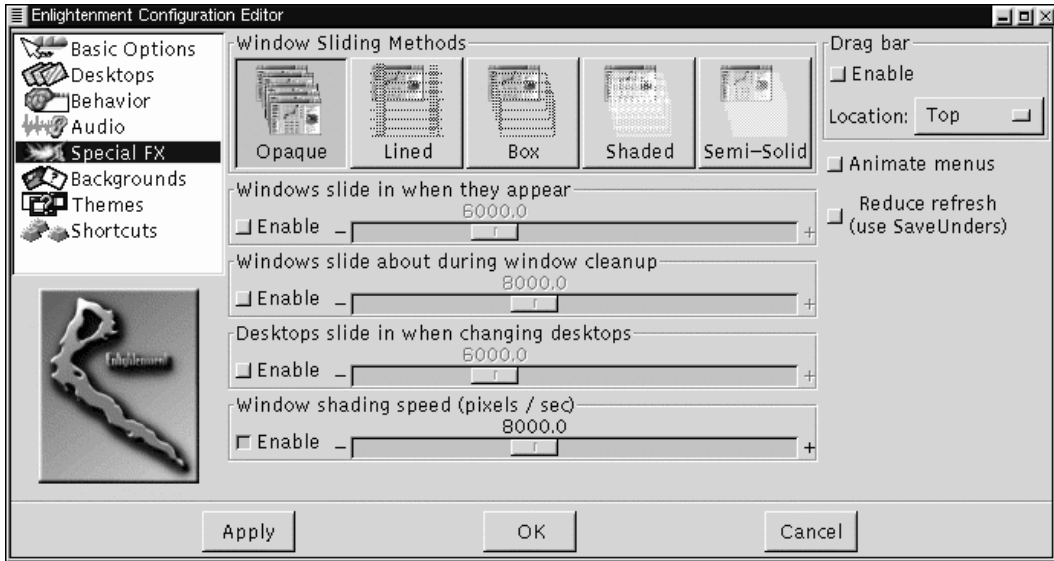
11.2.4 Audio

You can turn on Enlightenment's own sound events in this section.

11.2.5 Special FX

The Special FX Section on E-conf controls various special effects that E can use to control windows. If you have a slower machine you might consider disabling these features in this section.

Figure 11–4   Enlightenment Special Effects

- **Windows Sliding Methods** - This area controls what windows look like when they slide in on first display. You must have the slide in feature enabled for these to work. You can choose from Opaque, Lined, Box, Shaded, and Semi-Solid.

- **Windows slide in when they appear** - This is where you can turn on or off the special effect of windows sliding in. To turn it on or off press the small button beside the speed slide bar. The speed slide bar controls how fast the windows slide and is based on pixels per second.

  Windows slide about when cleanup in progress - This controls the special effect of windows sliding in when you select the **Cleanup Windows** menu item from the Enlightenment menus. To turn it on or off press the small button beside the speed slide bar. The speed slide bar controls how fast the windows slide about and is based on pixels per second. For more information about the **Cleanup Window** function see Section 11.3, *Enlightenment Menus*.

  Desktops slide in when changed - If you use multiple desktops this will make the desktops slide in when you switch to a new one. To turn it on or off press the

small button beside the speed slide bar. The speed slide bar controls how fast the desktop slides in and is based on pixels per second.

Window shading speed - If you choose to have this option you can shade a window by double clicking on its title bar. Shading will draw the body of the window up into the title bar so that you only see the title bar on the screen. Another double click will return the window to it's original state. The time slide bar next to the button controls how fast this behavior works. The speed is measured in pixels per second.

• Drag bar - Some Enlightenment themes will include a Drag bar which allows you to slide and peak into the next Desktop area or Desktop. If you wish not to have this on, make sure it is not selected. You can also decide where to place the Drag bar if you do have it enabled.

• **Animate Menus** - This will control the way Enlightenment menus are displayed. The menus will be animated if you enable this feature by pressing the small button.

• **Reduce refresh by using more memory** - This will place menus into memory to speed them up. Depending on your menus this can place a sizeable hit on your system memory.

11.2.6 Backgrounds

You can choose to let Enlightenment control your desktop background images. Setting the background is controlled in this area of the Enlightenment Configuration tool.

IMPORTANT

If you do choose to have Enlightenment set the desktop background you will not be able to use the GNOME Control Center's Background Properties Capplet. If you are new to GNOME or Enlightenment it is recommended that you do not use this feature.

Figure 11–5   Enlightenment Backgrounds
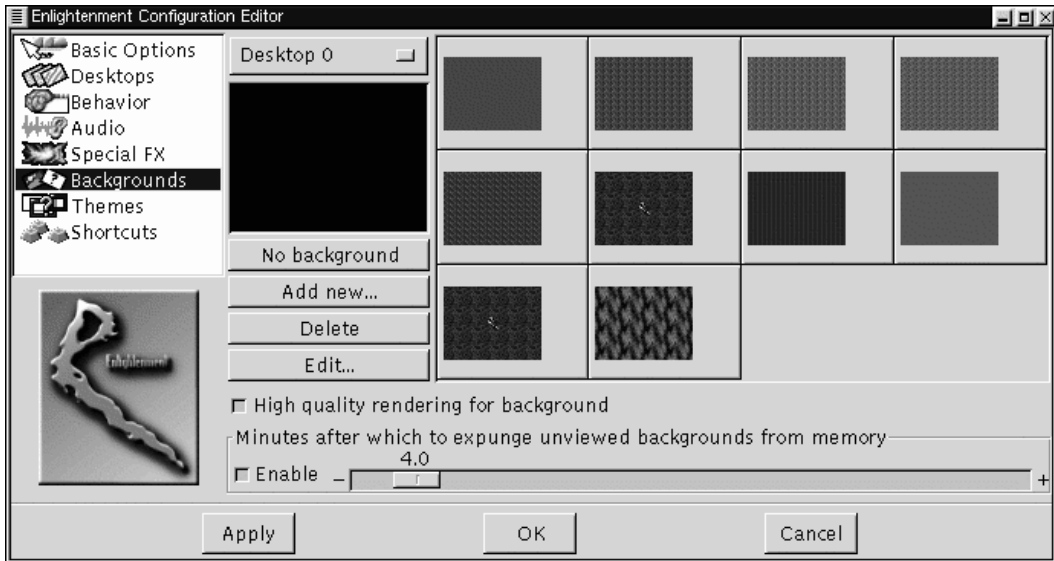
Setting the Background

The Enlightenment background selector acts as much as a file viewer for your back-grounds as it does a creator of them. You will notice first that there is a large area to preview the backgrounds. This area is not only a good place to preview your back-grounds but it acts as an area for "bookmarking" them.

To set a background for your desktop you must first choose which desktop you want to set the background for. The top button which reads **Desktop 0** will show you which desktop you are setting the background for. If you want to set it for another desktop you can press the button and select the proper desktop.

NOTE

Desktop 0 is the first desktop in any series of desktops. If you are using virtual areas on one desktop it is **Desktop 0** as well.

Once you have chosen the desktop you wish to set a background for you can select a background for that desktop. To select one simply locate the image or color in the collection on the right of the preview area or select the **No background** button.

If you do not like the selections in your collection you can create a new one. You will find more information about making a new background in the next section.

NOTE

If you select the **No Background** button for a desktop, another application or the GNOME Control Center's Background Properties Capplet can set the background for that particular desktop.

Creating a New Background

To create a new background you must first press the **Add new** button to start adding the new background to the current collection of backgrounds on the right. This will add a new box to the collection which will be blank at first.

You can start to edit the new image by selecting the new box you have just added and pressing the **Edit** button. This will bring up the **Edit Background** dialog which allows you to set the background image.

The **Edit Background** dialog contains three tabs which represent three different types of backgrounds you can have: **Solid Color**, **Gradient**, **Background Image**, and **Overlayed Logo**.

Figure 11–6   Solid Color Background
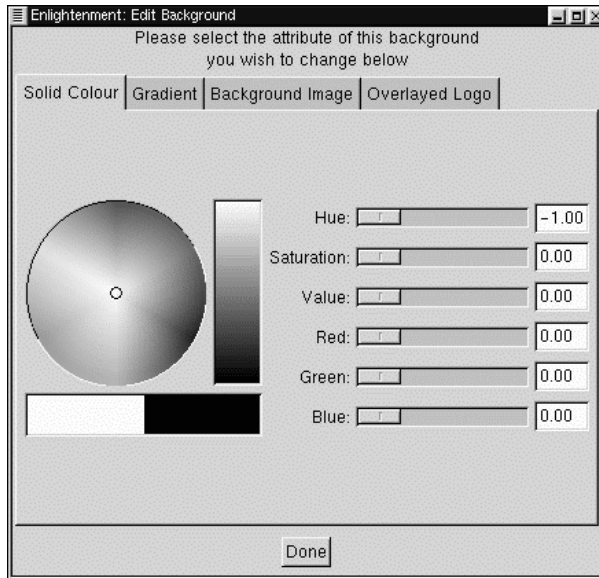
- **Solid Color** - This tab is simply a color selector and color properties slidebars which allow you to choose the solid color you wish to use for your background.

Figure 11–7   Gradient Color Background
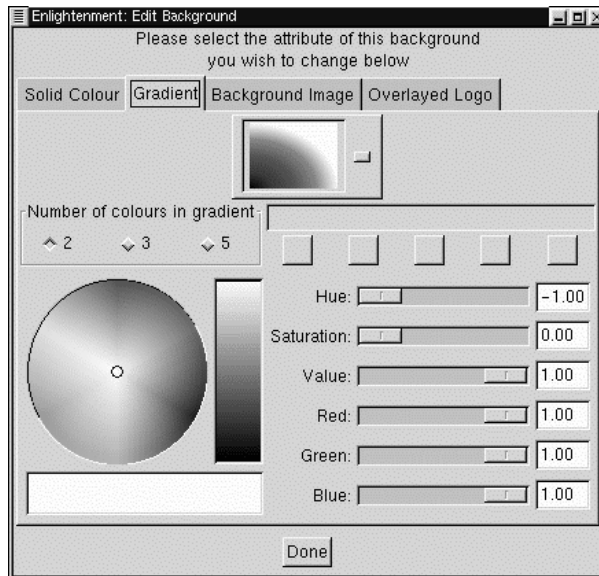
- **Gradient** - This tab combines the color selector and color properties slide bars with a choice of how many colors you would like in the gradient. You can select as many as five colors but you must specify the color for each selection.

Figure 11–8   Background Image
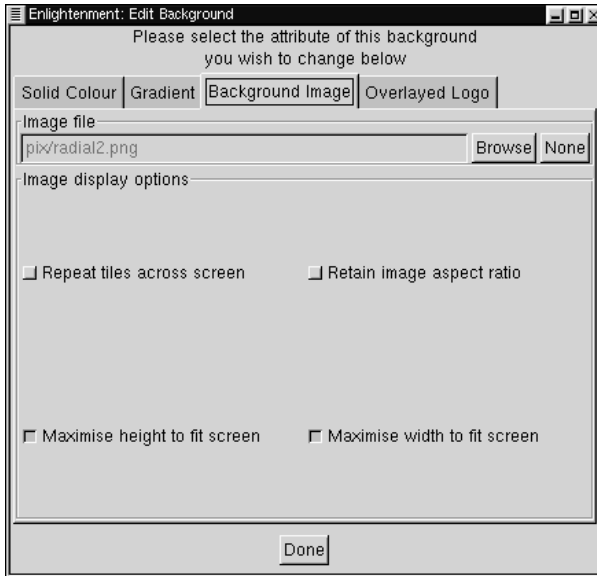
• **Background Image** - This tab allows you to select an image on your system to display as your background. You can press the **Browse** button to find the image and the **None** button to clear any selected image.

Once you have found the image you want you can decide whether the image should repeat as tiles across the screen, retain its own aspect ratio, maximize its height to fit the screen, or maximize its width to fit the screen.

Figure 11–9 Overlayed Logo
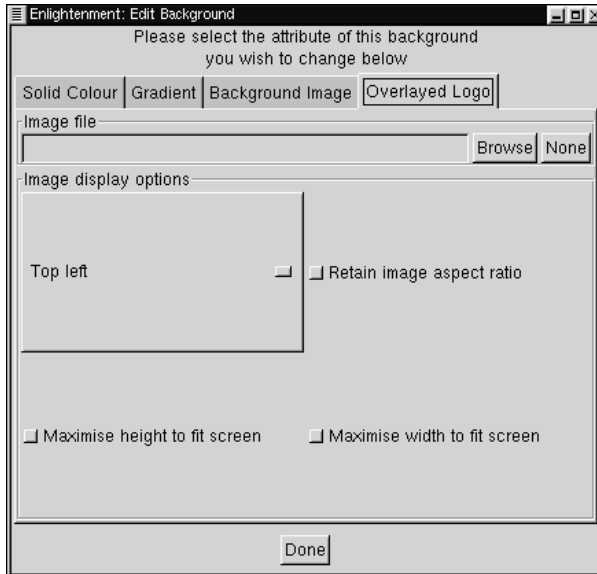
- **Overlayed Logo** - This tab will allow you to overlay an image on the background you have already set in the other tabs. You can search for the correct image with the **Browse** button and clear any selection with the **None** button.

  Once you have found the image you wish to overlay you can select where you would like the image to be displayed with the placement drop-down list. You can also decide if you want to maintain the image's original aspect ratio, maximize its height to fit the screen, or maximize its width to fit the screen.

When you have completed editing your image, press the **Done** button to return to E-conf.

11.2.7 Themes

Enlightenment is known for many things to the Linux community, not least of which is the sometimes fantastic and personalized look. This can be attributed to the one of the most fun features of Enlightenment, support of themes. These themes have been

created by Raster and many E users throughout the world. In fact, a Web site has been set up to be a repository of E themes, http://e.themes.org.

Figure 11–10   Enlightenment Themes
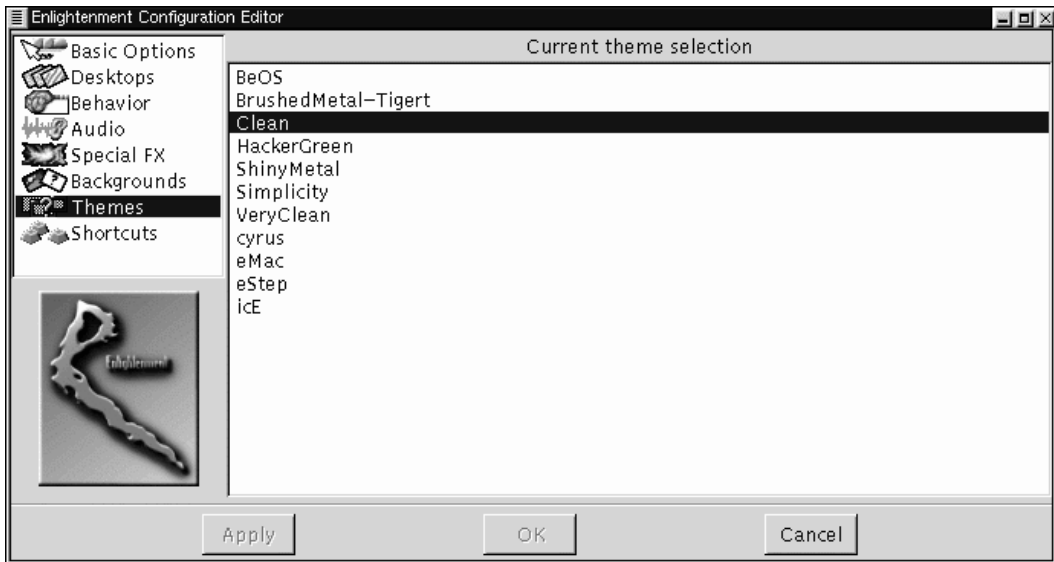
E themes must have an extension of .etheme to work. Some themes might be made for older version of Enlightenment and will not work with the newest release so make sure you are getting a compatible version.

Once you have found a theme you would like to have for E you must put it in your

```
~/.enlightenment/themes/
```

directory. All .etheme files in this directory will be available for selection in the **Themes** section of the Enlightenment Configuration tool.

To change themes select the theme you would like in the list of available themes and press the **Apply** button on the bottom of the E-conf.

11.2.8 Keyboard Shortcuts

The Keyboard Shortcuts section will allow you to change the keyboard shortcuts used for controlling windows, starting commands, and navigating between windows, areas, and desktops.

Figure 11–11   Enlightenment Keyboard Shortcuts

The Keyboard Shortcuts section is divided into two main areas, **List of keyboard shortcuts** and the **Edit current selected keyboard shortcut**.

In the **List of keyboard shortcuts** you can see which shortcuts have already been defined and select them to edit. You can also add new shortcuts by pressing the **New** button and delete them by pressing the **Delete** button.

The **Edit current selected keyboard shortcut** area is where you actually define what the keyboard shortcut will be. To use this area, you must select the shortcut you wish to edit, or select a newly added shortcut after you press the **New** button. Once a shortcut has been selected you can change the keys it uses or even the function it performs.

To change the keys used for a shortcut press the **Modifier** selection button and select which modifiers you wish to use. Modifiers are the keys [ALT], [CTRL], and [SHIFT]. You can select one modifier or a combination of modifiers to use with your shortcut.

At this point you will also need to change the key being used by pressing the **Change** button beside the **Key** definition box. Once you press the **Change** button you will receive a short message pop-up window which will tell you to press the desired key. Press the key and the definition will be changed.

To assign the new shortcut to an action simply scroll though the **Action to perform** list and select the action you wish to perform with this shortcut.

Some keystrokes you wish to use might have an action that needs defining. For example, if you wish to use a keyboard shortcut to change to a particular desktop you can select **Goto Desktop**. In the **Options for Action** you can define which desktop you want to go to by typing in the correct number.

11.3 Enlightenment Menus

Like all other window managers Enlightenment is designed to be used with or without a desktop environment like GNOME running on top of it. This means that there is some funcitonality in it which allows you to do things like start applications, move between desktops and more.

To this end, Enlightenment provides you with menus to launch applications and perform tasks. These menus may be used in one of two places that provide different functionality: the desktop and any window border.

11.3.1 Enlightenment Desktop Menus

To access the desktop menus, click with your middle mouse button (or both right and left mouse buttons if you have emulation enabled) on the desktop. This will bring up the Enlightenment menus.

The desktop menus will provide you with various applications that you might have installed on your system. There are some applications in the menu that you may not have installed which the creator of Enlightenment chose to put in the menus.

The GNOME main menu items are mirrored in the Enlightenment menus as well. Plus you can launch the Enlightenment Configuration tool or change themes on the fly.

Guide to Enlightenment Desktop Menus

IMPORTANT

These menus may be different if the creator of an Enlightenment theme changes them. If there are options not covered in this guide you should check the documentation associated with the theme you are using.

- **GNOME Apps** - This is a mirror of the GNOME Main Menu applications.

- **User Apps** - This is a mirror of any apps you have added to the User menu in the GNOME Main Menu.

- **Other Programs** - These are applications that the creator of Enlightenment or the creator of the theme you are using decides to include. You may or may not have these applications loaded on you machine.

- **Desktop** - This menu includes items that make changes to your desktops.

  - **Cleanup Desktop** - This will rearrange the windows on your desktop to a cleaner pattern by arranging them in a pre-determined order. This is very useful if you are having trouble finding an application due to the clutter on your desktop.

  - **Go to Next Destop** - This will take you to either the next desktop or the next desktop area you have, depending on which is next.

  - **Go to Previous Desktop** - This will take you to either the previous desktop or area, depending on which is next.

  - **FX Ripples** - If you have time on your hands and perhaps a good sense of humor, select this item and keep a close eye on the bottom of your screen.

- **Themes** - This will search the themes directory for any themes you have installed on your machine. You can select any theme you have available and change the theme on the fly.

- **Enlightenment Configuration** - This will launch the Enlightenment Configuration tool.

- **About Enlightenment** - This will launch a small window to give you more information on Enlightenment.

- **Help** - This will launch Enlightenment's Help System

- **Restart Enlightenment** - This will restart Enlightenment completely. If you are using GNOME on top of Enlightenment this is not a recommended option, instead, you should use the GNOME Logout.

11.3.2 Enlightenment Window Menus

The menus that are provided on the window borders can be launched one of two ways. If the creator of the theme you are using has included it you can use a menu button. You might find that this is not the case for certain themes, in which case you can use the second selection method.

The second way to get the window border menus is to press the [Alt] key and the right mouse button anywhere on the window border or in the window itself. This will work with the default keyboard shortcuts that ship with Enlightenment but you should check and make sure that you have not changed this keystrokes in the Enlightenment Configuration Tool.

The Enlightenment window border menus allow you to do specific things to the window itself.

Guide to Enlightenment Window Menus

IMPORTANT

These menus may be different if the creator of an Enlighten-
ment theme changes them. If there are options not covered
in this guide you should check the documentation associ-
ated with the theme you are using.

- **Close** - This will close the window.

- **Annihilate** - This will kill the application if you are having trouble closing the
  window.

- **Iconify** - This will iconify your window. If you are using the GNOME Pager it
  will simply minimize the window to the task list on the GNOME Pager. If you
  are not running the GNOME Pager it will minimize it to an icon based on the
  applciation.

- **Raise** - This will raise the window to the front of all other windows on your
  desktop.

- **Lower** - This will lower the window to the bottom of all of the windows on the
  desktop.

- **Shade/Unshade** - This will "shade" a window or unshade it if it has already been
  shaded. Shading a window will "draw" the window up into the window titlebar
  so that all you see is the titlebar itself.

- **Stick/Unstick** - This menu item allows you to make a window visible on all desk-
  tops and desktop areas.

- **Desktop** - The items in this menu allows you to move a window to another desktop
  area. The selections available are: **Move to Area Right**; **Move to Area Left**; **Move to Area
  Above**; **Move to Area Below.**

- **Window Size** - The items in this menu allow you to resize your window to certain ranges. The ranges are divided into Height, Width, and Size (both height and width). Within these ranges you can select how you wish to resize to the ranges.

12   Using the GNOME Panel

12.1 Introduction

The Panel is the heart of the GNOME interface and acts as a repository for all of your system applications, applets, and the Main Menu. The Panel is also designed to be highly configurable. The Panel gives you a place that always contains menus and applications as you want them to be.

12.2 The Basics

Using the GNOME Panel is very simple and will come easily to anyone who has used a graphical based operating system. You may add new panels, add applications to the panel, and add various applets. All of these functions and more will be described in this section.

12.2.1 Using the Main Menu

To start using any pre-loaded application click on the **Main Menu Button**. The Main Menu has the picture of the stylized foot and on first use is on the bottom left of the screen. You should release the mouse after pressing the **Main Menu** button so that you can take advantage of other mouse-activated features in the Main Menu such as right mouse clicks and drag and drop from the menu.

Figure 12–1   The Main Menu Button

The Main Menu is the starting point for all of the applications on your system. Later in this manual you will learn how to customize the Main Menu to suit your work environment, but for now you can use the menu that is established when you install GNOME. The Main Menu works like any other menu you might have used in other

graphical desktop environments. Simply press the **Main Menu** Button and select from
the menu that pops up from the button.

12.2.2 Hiding the Panel

At any time you may hide the GNOME Panel by pressing the **Hide Button**

Figure 12–2   The Hide Button

This will hide the GNOME Panel in the direction of the arrow on the **Hide Button**.
There are hide buttons on both sides of a Panel so you can hide it in either direction.

You may decide that you want the Panel to hide on its own when you are not using
it. This can be a helpful function if you are unable to run your system in a high
resolution. You can find out how to autohide the Panel in Chapter 15, *Configuring
the Panel*.

12.3 Moving and Adding Panels

Any Panel you have on your desktop can be moved by using the middle mouse button,
or by simultaneously pressing the left and right mouse buttons, to drag the panel to
the desired edge of your screen. If you do not have a middle mouse button and did
not configure your mouse to emulate a middle button you may also move a Panel by
changing its location in the Panel Configuration dialog. You can read more about this
in Chapter 15, *Configuring the Panel* of this documentation.

You may also add a new Panel to your desktop by selecting the **Add New Panel** from
the **Main Menu** | **Panel** menu. You will be given a choice of Edge or Corner Panels.
Both of these Panel types are described below.

•   Edge Panel - An Edge Panel is exactly like the main Panel that starts up with
      GNOME. By selecting this type of panel you may add a new Panel to another
      edge of your screen to give yourself more functionality.

- Corner Panel - The Corner Panel is a small Panel that will not stretch across the entire edge of the screen it is on. The Corner Panel will, however, stretch to the extent needed to display the icons and applets it contains. The hide buttons work just a bit differently with Corner Panels. The hide button closest to the edge of your screen will hide the Panel as usual but the other hide button will send the whole panel to the opposite edge. When the latter move is made it will not hide the Panel since it is changing the side of the screen it resides on. If you want to hide it you will have to press the hide button once again.

At any time you can change the current panel to the opposite type by selecting either the **Convert to edge panel** or the **Convert to corner panel** from the pop-up menu. The selection that is available depends on which type of panel you right mouse click on. The selection displayed will be the opposite of the current panel.

12.4 Adding Applications and Applets to the Panel

12.4.1 Adding Application Launchers

If you would like to add an Application Launcher (an icon that starts a particular application) to the Panel, right mouse click the Panel and select **Add New Launcher** from the **pop-up**

After selecting the **Add New Launcher** menu item you will see a dialog which will allow you to set the properties for the application launcher you wish to add.

Figure 12–3   The Create Launcher Dialog
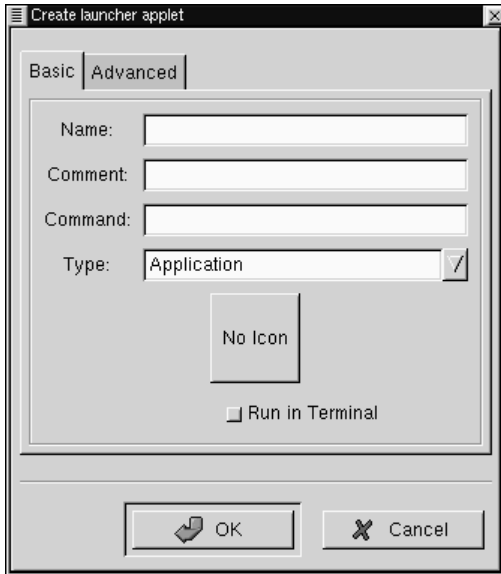
In the Create launcher applet dialog you may add a name for your launcher, a comment, the command line to launch the application, and define the application type. You may also press the icon button and choose an icon to represent the application from the icon picker. If no icon is chosen a default icon will be used.

Another, quicker method of adding an application launcher to the Panel is to go into the Main Menu and right mouse click on an application menu item. You will be given another menu selection which contains **Add this launcher to Panel**. If you select this menu item it will automatically add a launcher for that application to the Panel in which you invoked the Main Menu. At this point you may right mouse click on the launcher and select the **Properties** menu item to change any options for that launcher.

12.4.2 Grouping Items with Drawers

If you would like to group a subset of applications together you may use a Drawer. A Drawer is simply a small menu-like button that sits on your panel that groups application launchers together is one place. Once you have placed a Drawer on the Panel you may click on it to raise the menu of applications and click again to lower them.

Figure 12–4   An open Drawer

There are a couple of ways to place a Drawer on your Panel. First you may right mouse click on the Panel and select the **Add drawer** from the **pop-up** menu. Second, if you want a whole subset of menus from the Main Menu to become a drawer you may right mouse click on the title bar of that menu and select the **Add this as drawer to panel** from the **pop-up menu**.

You can add menus to your panel in the same way you add Drawers. Menus are very similar to Drawers except that they do not use large icons to represent application launchers, instead they use a style similar to the Main Menu, i.e. small icons and the

application name. You may add a menu by right mouse clicking on the title bar of a menu and select the **Add this as menu to panel** from the **pop-up menu**. You may also add system directories to the Panel as menus by dragging a directory out of the GNOME File Manager and dropping it on the Panel.

12.4.3 Adding Applets

There are many applets which you can add to the Panel as well. Applets are small applications which can perform tasks within the Panel itself. There are many applets you can add to the Panel and these are covered in more detail in Chapter 19, *Panel Applets* of this manual. As an example of how to add an applet you can add another clock applet to your Panel. To add the clock applet to your Panel:

- Right mouse click on the Panel
- From the **pop-up** menu choose the **Add new applet** menu item.
- This will bring up more levels of **pop-up** menus.
- Choose the **Utility** | **Clock** menu item.
- The Clock will be added to your Panel.

Figure 12–5   The Clock Applet
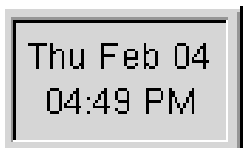
To remove the clock applet you added you can right mouse click on the clock and select **Remove from panel** from the **pop-up** menu.

12.5 Running Applications

There are many ways to launch applications within GNOME. Remember that one of GNOME's strong points is that it allows you to start and control applications with an easy-to-use graphical interface. There are many ways to start the applications you wish to use:

As you saw in the previous section you may use the Main Menu to find applications which have been pre-loaded or you provided using the Menu Editor. You can read more about the Menu Editor in the Chapter 16, *Editing the Main Menu*.

You may also add application launchers from the Panel, which was covered in Section 12.4.1, *Adding Application Launchers*.

If you are using the GNOME File Manager, you may double click on any executable file and it will run.

You may use the GNOME Run program to launch any application. To use it select the **Run program** menu item from the **Main Menu**. This will launch a simple dialog that allows you to type in the command for launching the application. For example, if you wish to start the Emacs editor and it is not in a menu or on your panel you may select the GNOME Run Program and type emacs in the text box on the dialog.

12.6 Logging Out of GNOME

GNOME has a couple of helpful methods for Logging out. You may either use the **Logout** menu item or the **Logout** button.

The **Logout** menu item is the bottom menu item in the **Main Menu**. Simply select the **Logout** menu item and you will be prompted on whether you would like to log out or not. Select **Yes** and your GNOME session will end.

If you would like, you may add the **Logout** button to the Panel. This is just another method for logging out. It provides no functionality beyond providing a pretty button. To add the **Logout** button to the Panel, right mouse click on the Panel and select the **Add logout button** menu item from the **pop-up** menu. Once the button is there you may press it to log out.

NOTE

If you are running a window manager that is GNOME compliant, the logout feature will quit the window manager as well as GNOME. If you are running a non-compliant window manager you will have to end that window manager yourself.

The Logout dialog will display when you log out of GNOME. This dialog will ask you whether or not you really want to log out. It also provides you with different methods of quitting GNOME.

You have three choices on how to quit GNOME. You may **Logout** which will simply take you to a terminal; **Halt**, which will shut down the whole system; or **Reboot**, which will reboot the whole system.

If you do not want to log out you may press the **No** button and you will be returned to your GNOME session.

Figure 12–6   The Logout Dialog

Within the Logout Dialog there is one option you may choose before you leave GNOME.

If you would like to save your current setup, you may select the **Save current setup** checkbox. This will save those programs you have open, and the configuration of your Panel.

13   The GNOME Desktop

13.1 Introduction

The GNOME Desktop provides the functionality of any traditional operating system desktop. You can drag files, programs and directory folders to the desktop; you can also drag those items back into GNOME-compliant applications.

IMPORTANT

The GNOME Desktop is actually provided by a backend process in the GNOME File Manager. If, for any reason, that backend process has stopped running you may start the GNOME File Manager again and your desktop will be restored. Even if you do have to do this you do not need to keep the GNOME File Manager window open to enable the desktop.

13.2 Using the Desktop

Using the Desktop is as simple as dragging items you wish to use routinely to the desktop. The default desktop will include a folder of your home directory (/home/[user name]). By default the GNOME File Manager window will also appear for you to access other areas of your system.

To utilize drag and drop you need to be using either a GNOME compliant application or a Motif application. GNOME is compliant with Motif drag and drop so you will find it works with many applications you already have installed.

All items that are stored on your desktop are located in the following directory:

```
$/home/[user name]/.gnome-desktop/
```

This is helpful to remember when you want your desktop to contain an item for which you can not utilize drag and drop.

Once you have started GNOME you will notice that any drives you have connected to your system will be shown on your desktop with the appropriate icons. You may mount and access these drives utilizing these icons.

IMPORTANT

You must have permission to mount the device shown on your desktop before you can utilize these icons. If you do not have mount permission someone with root access such as your system administrator can give it as follows.

Giving mount access to ordinary users can be done quite easily if you have *linuxconf* installed on your machine. Just select the drive you want to access in the **Access local drive** section. In the **Options** tab select the **User Mountable** option. Your drive will now be mountable by users.

If *linuxconf* is not available someone with root access must edit the */etc/fstab* file to include user access. This is done by adding the user attribute to the drive. For Example:

If your fstab file contains a line like this:

```
/dev/cdrom /mnt/cdrom iso9660 exec,dev,ro,noauto 0 0
```

Add "user" to the fourth column:

```
/dev/cdrom /mnt/cdrom iso9660 user,exec,dev,ro,noauto 0 0
```

Depending on your system and work environment, there could be some security risks in permitting users to mount disks. Please consult your system administrator before taking this action.

Once you have permission to mount a drive you may right mouse click on the drive icon on your desktop. This will bring up a small pop-up menu.

You may select **Mount device** to mount it and **Eject device** to eject it. Once it is mounted you may either double-click it or choose **Open** from the pop-up menu to open and GNOME File Manager window viewing the contents of the device.

If you are missing any drives that might have been added to your machine you may right mouse click on an empty space on the desktop and choose the **Recreate Desktop Shortcuts** menu item from the **pop-up** menu.

13.3 Desktop Areas

Desktop areas allow you to keep a well organized system when you have many tasks to perform at one time. Just like adding a new desk when you have too much material to fit on one, desktop areas allow you to move to another area to launch more programs.

GNOME is aware of desktop areas even though they are controlled by another software program called the 'window manager'. You can set the number of desktop areas within the configuration of the window manager you are using. If you are using the default window manager or your window manager has a graphical configuration tool you may be able to launch it from the Window Manager Capplet. You may read more about this Capplet in Section 17.2.4, *Window Manager Capplet*

IMPORTANT

Most window managers will give you the option of having multiple desktops, which are different from desktop areas. Desktop areas are virtual extensions of one desktop whereas multiple desktops are actually separate.

The default setup of GNOME is to use desktop areas with only one desktop. The reason for this is with some applications, such as those which use Motif, users can experience problems with some drag and drop functionality across desktops.

13.4 Other Desktop Menus

There are a few desktop menus you may use in GNOME. These menus are accessed by making a right-button mouse click on any clean space on the desktop. This will bring up a **pop-up** menu which contains a few items:

- **New | Terminal** - This launches a new GNOME Terminal window that will automatically navigate to the ~/.gnome-desktop directory.

- **New | Directory** - This creates a new directory on your desktop into which you can place files: a convenient way to clean up your desktop.

- **New | Launcher** - This allows you to place a new application launcher on the desktop. When selected, this menu item will launch an **Application Launcher** dialog that allows you to specify which application and its properties.

- **New | [application]** - Some applications may put items for you in the **New** menu. For example, the Gnumeric spreadsheet will put a **New Gnumeric Spreadsheet** menu item in the **New** menu so you can start up a new spreadsheet easily.

- **Arrange Icons** - This automatically arranges your desktop icons.

- **Create New Window** - This launches a new GNOME File Manager window displaying your Home directory.

- **Recreate Desktop Shortcuts** - This rescans the mountable devices on your machine and display an icon for any new device it might find. It also reloads any desktop items your distribution places on the desktop by default.

- **Rescan Desktop** - This rescans the files in your ~/.gnome-desktop directory. Any new files found will be placed on the desktop as icons.

14   The GNOME File Manager

14.1 Introduction

GNOME includes a file manager that allows you to manipulate the files on your system in a comfortable, powerful, graphical environment. This File Manager is known as GMC for GNU Midnight Commander. GMC is based on the Midnight Commander file manager which can be run in a terminal.

Midnight Commander has long been known for its power and ease of use. GMC has taken the power and ease of use from MC and added the GNOME graphical frontend.

As mentioned in the previous chapter, the GNOME File Manager provides the desktop functionality for GNOME. The GNOME File Manager also provides a place to manipulate files on your system by using the GNOME File Manager window.

Figure 14–1   GNU Midnight Commander - The GNOME File Manager
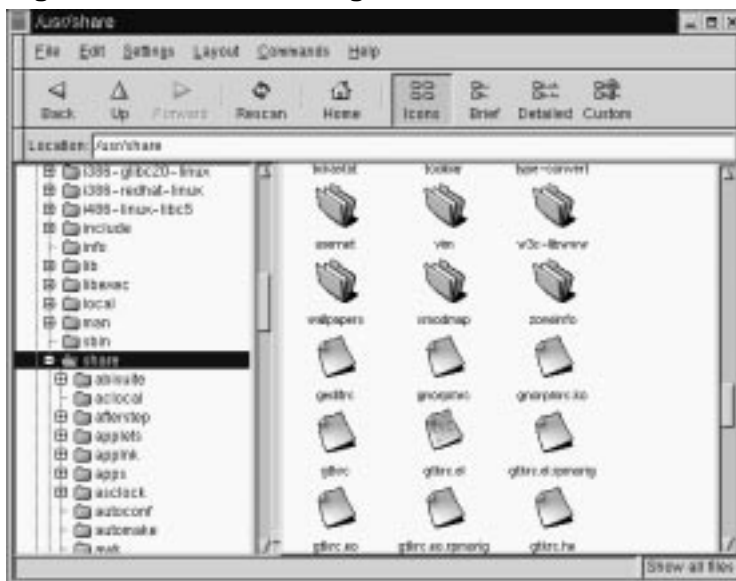
There are two main windows within the GNOME File Manager. On the left is the tree view, which represents all of the directories on your system by their hierarchical position. On the right is the directory window, which will show you the contents of the directory which you have selected in the tree view.

To select a directory in the tree view simply use a single mouse click. This will change the main directory view, showing the files in the directory you have chosen. If there is a plus sign to the left of a directory name in the tree window, you may click on it to show subdirectories. At this point the plus sign will have changed to a minus sign. If you click on the minus sign, the subdirectories once again become hidden in the tree view and the plus sign reappears.

The directory window has a few viewing options you may wish to take advantage of: the icon view and and three list views, brief, detailed, and custom.

The Icon view is the default view and will display large icons for each file. The Brief view shows the files and directories in a list. The Detailed view expands this list to include information about the files. The Custom view is a list view that allows you to select the information you want to view about files.

In the Brief, Detailed, and Custom views if you click on one of the information titles on the top of the window it will sort the files according to that information. For example, if you want to find the largest files in the directory you can click on the Size title and the files will be sorted by size from largest to smallest. One more click on the Size title will change the sorting from smallest to largest.

14.2 Moving Around the GNOME File Manager

Above the tree and main file windows in the GNOME File Manager there is a Location text box.

Within this text box you can type the path location of the file in which you would like to view in the main window.

The GNOME File Manager is also equipped to view FTP sites. To view an FTP site you will need to be connected to the Internet either through a dial up account or a network. Type the FTP address in Location text box using and the GNOME File Manager will attempt to connect to the site.

Make sure you type in FTP addresses in the following manner:

```
ftp://[site address]
```

Above to the Location text box are navigation buttons you might wish to use. These buttons with the arrow icons allow you to move within the directories adjacent to the one you are in. The left button will take you to the previous directory you were in, the middle button will take you up one level in the directory hierarchy, and the right button will take you to the next directory if you have just moved backwards.


TIP

You can open a secondary window by using your middle mouse button on any directory. Press the middle mouse button (or both left and right mouse buttons if you have emulation) on a directory in the main window and a new window will open showing the contents of the directory you clicked on.

14.3 Selecting Files

Selecting files is done with your mouse by clicking on the file or files in the GNOME File Manager. The file that is selected will then highlight to show you that it has been selected.

There are a couple of ways to select more than one file. One way is to use the "rubber band" select by clicking and dragging the mouse cursor around several files. This action will produce a small dotted line, the "rubber band", to show you the area in which files will be selected.

Figure 14–2   "Rubber Band Select"
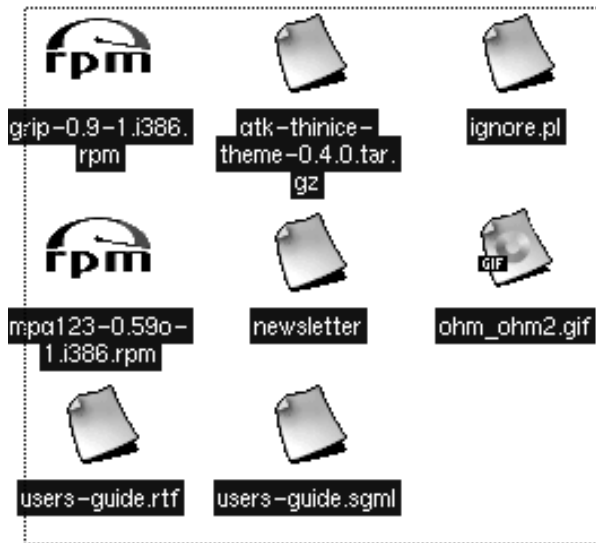
If you wish to be more selective about the files you are choosing, or the files you need do not reside next to one another, you may use the [CTRL] key to keep the files you have selected while you are selecting more. This works by selecting a file, pressing and holding the [CTRL] key and selecting another file. While the [CTRL] key is pressed you will be able to add to the "list" of files that are selected. Once you have selected multiple files by either method you may copy or move the files.

You may also select all files in a directory by selecting the **Select All** menu item from the **Edit** menu.

You may also filter your selection by using the **Select Files** menu item in the **Edit** menu. Using **Select Files** will display a simple dialog which will allow you to type in criteria for your selection. In this field, the symbol [*] is interpreted as a wildcard, E.g., it matches any string. For example, if you would like to select all files in the directory that start with the letter [D], you can type [D*] (note that filenames in Unix are case-sensitive).

14.4 Copying and Moving Files

The default action for dragging and dropping files between the GNOME File Manager and the desktop is to move them. But you can also use drag and drop to copy a file by pressing the [SHIFT] key while dragging the file(s). You can also drag a file within the file manager to another directory. Select the file in the main window and drag it to the desired directory in the tree window. Note, however, that when dragging to the tree window the default action is to copy the file. As you drag the file to the tree window, you will see a plus sign appear to remind you that the file will be copied. You can also use drag and drop to copy directories the same way as you did for files. Again note the plus sign signifying that the directory will be copied not moved

You may also bring up a menu which lets you decide what action to perform with a drag by using the middle mouse button to drag a file or pressing the [ALT] Key while dragging a file. Once you release the drag you will get a pop-up menu which contains the options Copy, Move, Link, and Cancel Drag.

Figure 14–3   The Move Dialog
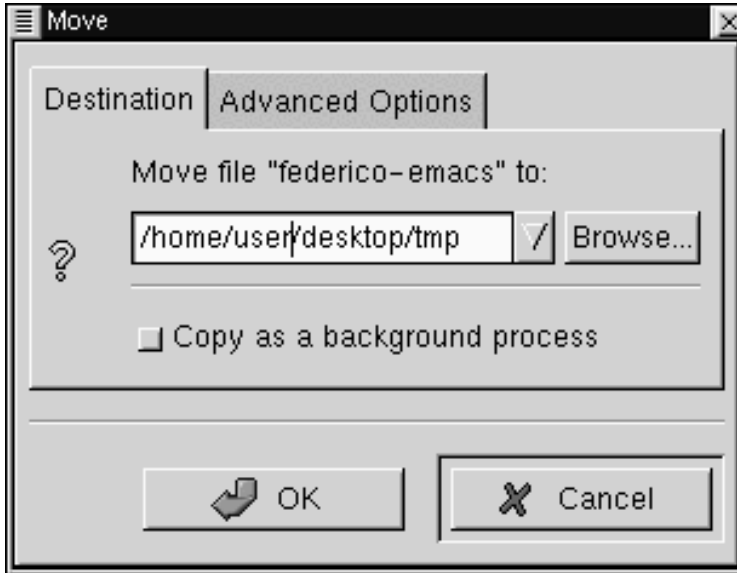
To use the move dialog you simply type in the path where you wish to move the file. If you want to rename the file you may type the new name of the file in the path string. The Copy dialog looks and works exactly the same way as the Move/rename dialog.

If the file you are moving is a symbolic link — that is, a virtual link to where the file actually resides (called a shortcut in some other operating systems) — you may select the **Advanced Options** tab and select **Preserve symlinks**. Selecting this will make sure the link is preserved despite the move.

It is recommended that you use this method of moving a symbolic link.

14.5 Renaming Files

Renaming files in the GNOME File Manager window or on the desktop can be achieved in two ways.

One method of renaming a file is to right mouse click on the file and choose the **Properties** menu item from the **pop-up** menu. In the Filename text box you may type in the new name as you wish it to appear.

**Figure 14–4   File Properties**
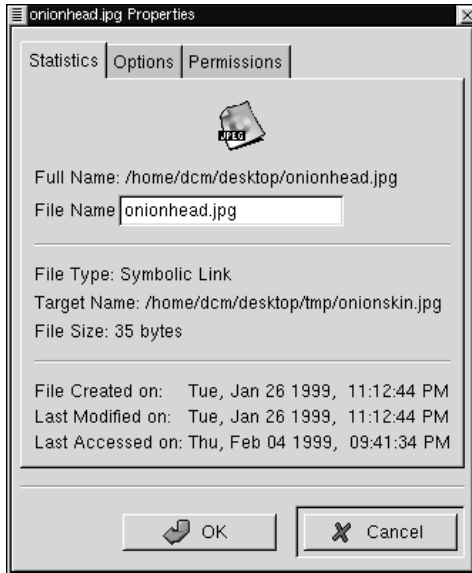
A shortcut method to rename a file is to slowly double click on the filename in the icon view or on the desktop. Make sure the double-click is slow so you do not launch the file but you simply highlight and then click again. At this point the name of the file will enter into the editing mode, your mouse cursor will change to an editing line, and you can type in the new name.

14.6 Launching Applications From the GNOME File Manager

The GNOME File Manager allows you to launch applications from the main window by simply double clicking your mouse on a file which has an application associated with it. You can change the way file types are handled by reading Section 17.4, *GNOME Mime Types*

If the file does not have an associated application you can right mouse click on the file and select **Open with** from the **pop-up** menu. This will bring up a dialog which that you to define the application which will launch the file. For example, if you want to edit a file names test.txt with Emacs (a popular text editor) you can right mouse click on test.txt and choose **Open with**. When you get the Open with dialog box simply type in emacs in front of the test.txt file name. When you press **OK** Emacs will open the file.

Figure 14–5   The Open With Dialog
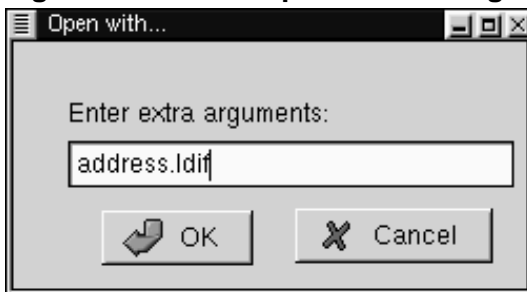

14.7 File Properties and Actions

For any file in the main file display you may right mouse click on it and choose a variety of properties and actions for it from the **pop-up** menu:

- **Open** - This will open the file with the default application associated with it. You may read more about editing these associations in Section 17.4, *GNOME Mime Types*.

- **Open With** - You may open a file with any application using this menu item. You may read more about this in Section 14.6, *Launching Applications From the GNOME File Manager*.

- **View** - This will view the file with a basic text viewer.

- **Edit** - This will launch an editor to edit the file. The editor launched is determined by the application associated with that file type. You may read more about editing this association in Section 17.4, *GNOME Mime Types*.

- **Copy** - This will copy the file to the clipboard so that it can be pasted elsewhere.

- **Delete** - This will delete the file.

- **Move** - This will bring up the **Move** dialog, which will allow you to move the file. You can read more about this dialog and moving files in Section 14.4, *Copying and Moving Files*.

- **Symlink** - Symlink allows you to create a symbolic link of the file you have selected in another directory. When you select Symlink you will see a dialog which allows you to type in the directory and filename that you wish to be the symbolic link to the file. This link will be a "shortcut" to the selected file.

- **Properties** - The Properties menu item will launch the Properties dialog. The Properties dialog allows you to edit and view the properties for the selected file.

  The Properties dialog consists of three tabs, Statistics, Options, and Permissions.

Figure 14–6   File Properties
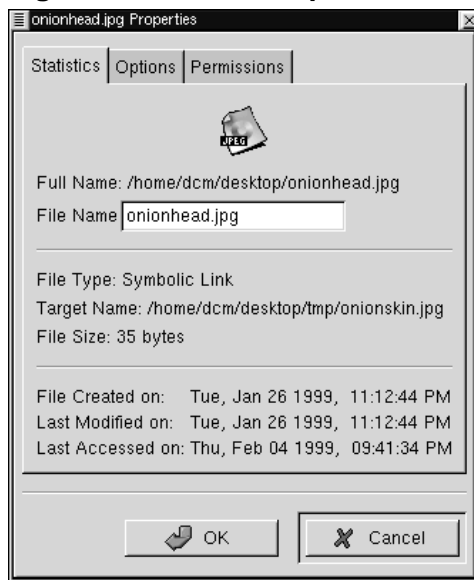
- **Statistics** - This tab will show you the file information including the name, type, size, and history. You may change the name of the file in the **File Name** text box.

- **Options** - This tab will allow you to change the action options for the file. You can define how to open, view, and edit the file. If you need to open it in a terminal window you may select the **Needs terminal to run** checkbox.

NOTE

If you bring up the Properties dialog from an icon on the desktop you will be able to change the icon for that file in the **Options** tab.

- **Permissions** - This tab allows you to change the permissions and ownership of a file if you have access to do so. You may select Read, Write, and Exec permissions for the User, Group, and Others. You may also set the UID, GID and Sticky as well as define who owns the file.

# 14.8 Changing Your Preferences in The File Manager

There are many settings you can configure for the GNOME File Manager.

These settings may be accessed from the **GMC Preferences** dialog. You may launch this dialog by selecting the **Preferences** menu item from the **Settings** menu.

The **GMC Preferences** dialog is divided into Five major sections: **File Display**, **Confirmation**, **VFS**, **Caching** , **Desktop**, and **Custom View**.

14.8.1 File Display

Figure 14–7   File Display Preferences
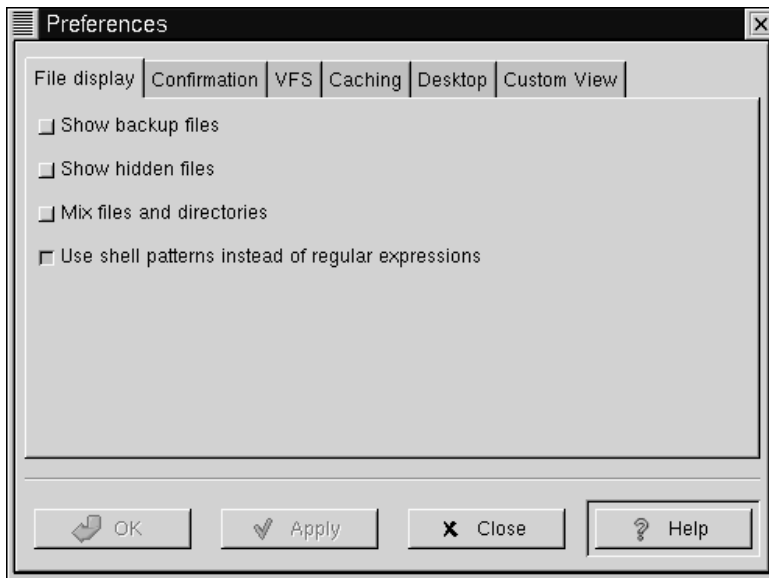
- The **File Display** tab allows you to change the way files are displayed in GMC.

  - **Show backup files** - This will show any backup file that might be on your system.

  - **Show hidden files** - This will show all "dot files" or files that begin with a dot. This files typically include configuration files and directories.

  - **Mix files and directories** - This option will display files and directories in the order you sort them instead of always having directories shown above files.

  - **Use shell patterns instead of regular expressions** - This option is for advanced users only. If you are unfamiliar with Regular Expressions you should not use this option. If you are familiar with how to create regular expressions you may select this option to use them in your sorts and filters.

14.8.2 Confirmation

Figure 14–8   Confirmation Preferences
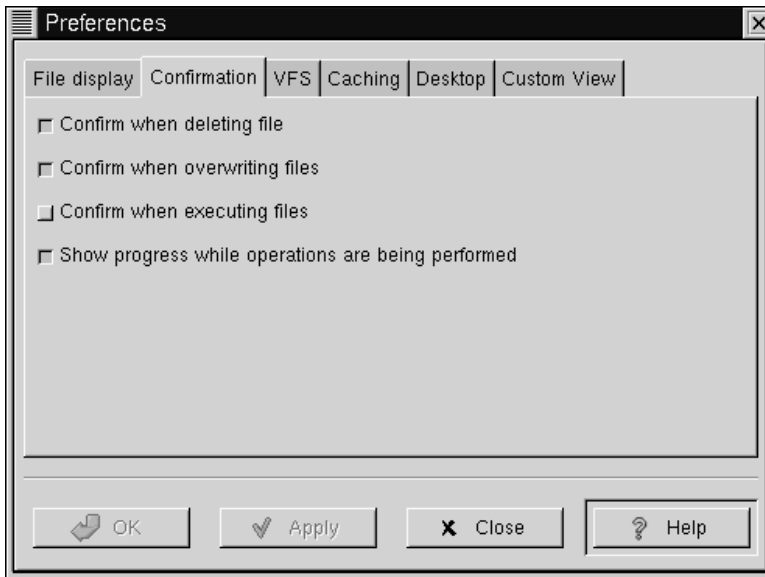
- This tab allows you to change which functions ask for your confirmation before continuing.

    - **Confirm when deleting file** - This will bring up a confirmation screen before deleting a file.

    - **Confirm when overwriting files** - This will bring up a confirmation screen before overwriting a file.

    - **Confirm when executing files** - This will bring up a confirmation screen before executing a file.

    - **Show progress while operations are being performed** - This will bring up a progress bar while certain operations are being performed such as copying, moving, deleting, etc.

14.8.3 VFS

Figure 14–9   VFS Preferences
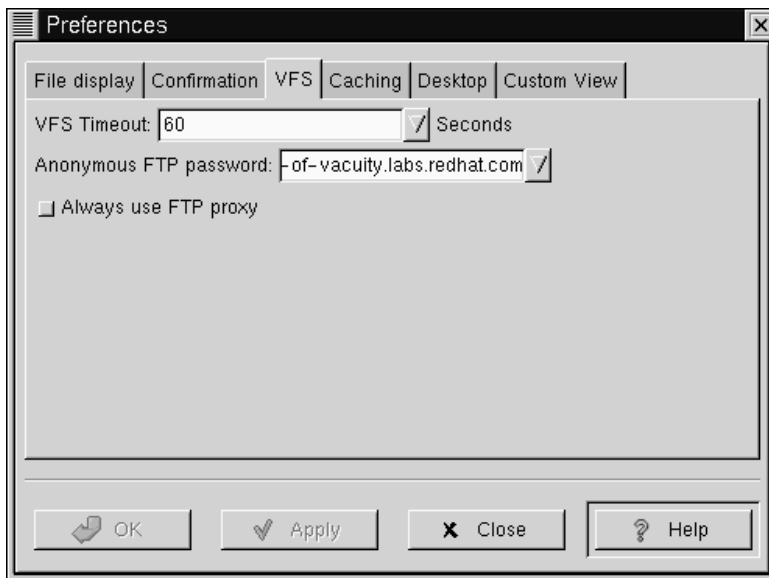
- This tab allows you to configure the options for your Virtual File System.

  The Virtual File System allows you to manipulate files that are not located on your local file system. There are different version of the VFS including ftpfs and tarfs. The ftpfs allows you to work on FTP sites while the tarfs gives you access inside .tar files.


NOTE

The .tar file is the standard UNIX archive format

- **VFS Timeout** - This will determine how long you will be connected to any VFS without activity. The timeout is measured in seconds.

- • **Anonymous FTP password** - This allows you to set a password for logging into anonymous FTP sites. Usually you will want to make this your email address.

- • **Always use FTP proxy** - If you need to use a proxy to connect to FTP sites you will want to enable this.

14.8.4 Caching

Figure 14–10   Caching Preferences
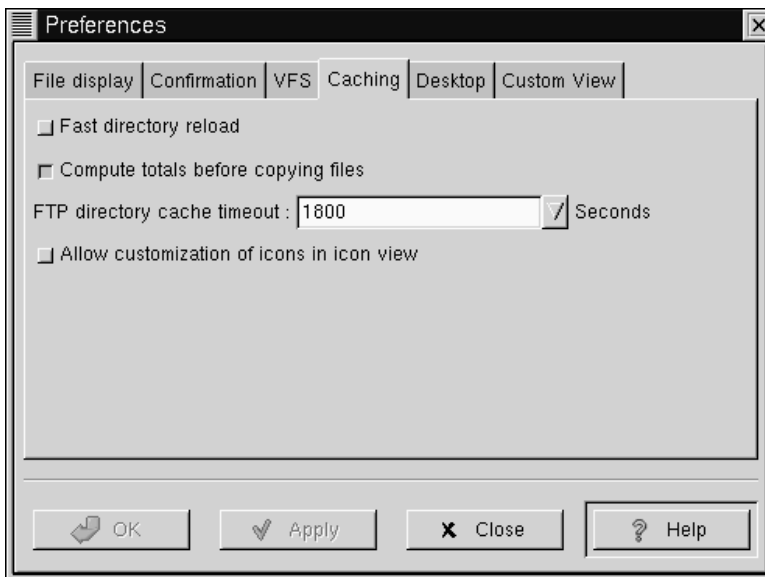
- • This tab allows you to configure items that will enhance the speed of GMC by using caching.

  - • **Fast directory reload** - This option will store directory information in cache so that it can load faster.

IMPORTANT

If you enable the **Fast directory reload** you may expe-
rience problems with not seeing new files that have
been added to directories.

- **Compute totals before copying files** - This will make GMC determine the number
  of files you are copying before it performs the task so that it can give you
  information about the process as its happening.

- **FTP directory cache timeout** - This option will keep recently visited FTP site
  information in cache for the amount of time you specify. The time is measured
  in seconds.

- **Allow customization of icons in icon view** - This will allow you to change the icons
  in the icon view by right mouse clicking on them and selecting the **Properties**
  menu item. You can always change the icons of items that are on your desktop
  but you can only change the icons in the icon view with this option turned on.

IMPORTANT

Turning on the **Allow customization of icons in icon view**
function may result in a slower system.

14.8.5 Desktop

As you learned before, the GNOME desktop is controlled by the GNOME File Man-
ager. The Desktop Dialog allows you to change the settings of the desktop.

Figure 14–11   The Desktop Dialog
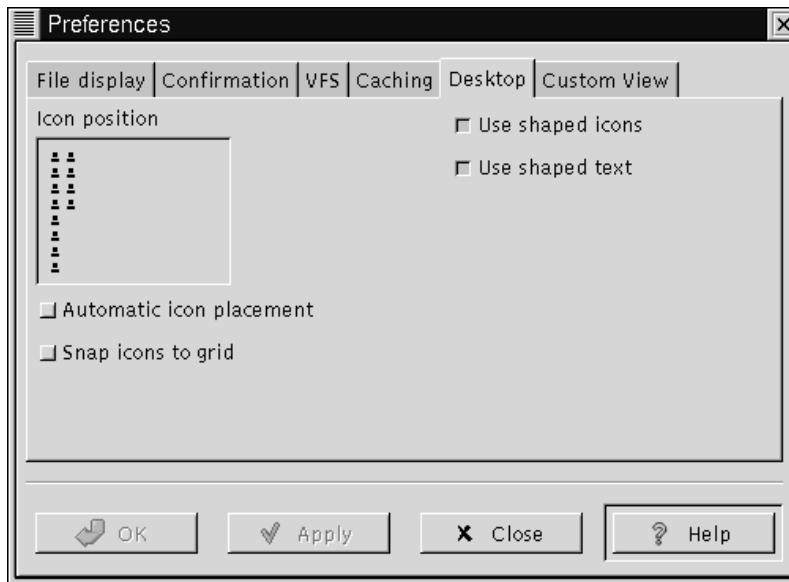
Icon placement - This small window allows you to determine where your desktop icons will be placed by default. To change the location of the icons simply click your mouse button in the window near the edge you would like the icons to appear. The small sample icons will move with your mouse clicks to show you where the default placement will be.

Use shaped icons - This option is turned on by default. If for some reason you have and older display, a strange X server setup, or any other graphic display problem, you may select this to make the desktop icons non-transparent.

Use shaped text - This option is turned on by default. Like the shaped icons above you may choose to have non-transparent text labels with your icons.

Auto icon placement - This will automatically arrange icons on your desktop for you. If you would rather place the icons manually you should turn this off.

Snap icons to grid - This will keep all of your icons in line when you move them to keep a clean desktop. It creates a grid that the icons will automatically move to when you release it after dragging.

14.8.6 Custom View

The Custom View dialog allows you to set the way you would like the Custom View to look in the GNOME File Manager.

Figure 14–12   The Custom View Dialog
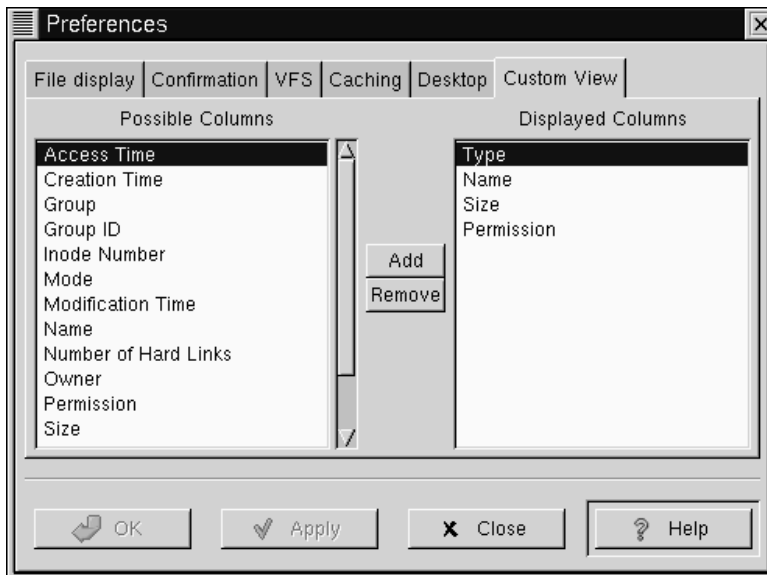
There are two main columns in the **Custom View**, the **Possible Columns** and the **Displayed Columns**.

The **Possible Columns** list includes all of the information types that can be displayed. If you would like to include one, you may press the **Add** button and it will be added to the **Displayed Columns** list.

The **Displayed Columns** list shows you the current information types that will be included in your Custom View. If you wish to remove any you may press the **Remove** button.

If you would like to re-arrange the items in the **Displayed Columns** list you may drag them to the desired location with your mouse.

14.9 Menu Guide to the GNOME File Manager

In this section each menu item in the GNOME File Manager will be described.

Tip

All menus in the GNOME File Manager are enabled with the tear-away feature. If you would like any menu to "float" on the desktop in its own window simply select the perforation line at the top of the menu.

- **File** - The File Menu contains items associated with files and higher level activity.

    - **Create New Window** - This will open a new File Manager Window.

    - **New | Terminal** - This launches a new GNOME Terminal window that will automatically navigate to the directory displayed in the main window view.

    - **New | Directory** - This creates a new directory in the directory displayed in the main window view.

    - **New | [application]** - Some applications may put items for you in the **New** menu. For example, the Gnumeric spreadsheet will put a **New Gnumeric Spreadsheet** menu item in the **New** menu so you can a new spreadsheet in the directory diplayed in the main window view.

    - **Open** - This will open the file you have selected with your mouse cursor.

    - **Copy** - This will launch the Copy dialog to enable you to copy the selected file to a destination you choose.

    - **Delete** - This will delete the file you have selected.

- • **Move** - This will launch the Move/rename dialog to enable you to rename the selected file or move it to a destination you choose.

- • **Show directory sizes** - This option will determine the size of your directories and display it in the Detail or Custom views. Depending on the size of your directories, this could take a while to display and may tax your system when viewing large directories.

- • **Close this Window** - This will close the GNOME File Manager window you are currently running.

- • **Edit** - The Edit menu contains items that are associated with editing and selecting files.

  - • **Select All** - This will select all of the files in the main window.

  - • **Select Files** - This will launch a simple dialog box which will allow you to specify a selection pattern. For Example, if you wish to select all files beginning with the letter D you can simple type D* and press OK.

  - • **Invert Selection** - This will invert the current selection. For example, if you have 7 files selected in a directory with 10 files, selecting Invert Selection will select the three remaining files and unselect the 7 files that were previously selected.

  - • **Search** - This toggles on the "completion" search for the Flie Manager. Once it is turned on the status bar at the bottom of the window displays the word "Search:" and allows you to start typing. If you are searching for a file named "GNU" in the directory displayed in the main window view typing "G" will move the selection to the first file or directory beginning with "G". Typing "GN" move the selection to the first file beginning with "GN", and so on. You may also toggle this feature on by pressing [CTRL]-[S].To end the search you may select the file or directory or simply navigate away from the file with the arrow keys or your mouse.

  - • **Rescan Directory** - This will rescan the current directory in case the files in the directory have changed.

- • **Settings** - This contains the Preferences menu item.

- **Preferences** - This will launch the Preferences dialog so that you can customize the GNOME File Manager.

- **Layout** - This menu contains items that are associated with the layout in GNOME File Manager.

    - **Sort By** - This will launch a dialog box which will allow you to select how you wish the files to be sorted in the current directory. You may sort by one of the following:

        Name

        File Type

        Size

        Time Last Accessed

        Time Last Modified

        Time Last Changed

    - **Filter View** - This will bring up a simple dialog which allows you to view only files with names matching a pattern. For example, if you wish to only view files in the directory that start with the letter D you can type D* into the dialog and press OK. If you wish to see all files again you will need to launch the Filter View dialog again and clear the dialog or type in a *.

    - **Icon View** - Selecting this will display the files in the main window as large icons.

    - **Partial View** - Selecting this will display the files in the main window as a list with only file names.

    - **Full View** - Selecting this will display the files in the main window as a list with all file information.

    - **Custom View** - The **Custom View** menu item switches the main window to the Custom View which is a list view displaying the information about your files that you specify. To customize the Custom view you will need to use the **Custom View Editor** in the **Preferences** dialog. You may read more about the

Preferences dialog in Section 14.8, *Changing Your Preferences in The File Manager*.

- **Commands** - The Command menu contains items that are commands to run on files in the GNOME File Manager.

  - **Find File** - Find File brings up a dialog which allows you to search for particular files on your system.

  - **Edit mime types** - This option will launch the GNOME Control Center Capplet that allows you to edit Mime Types for GNOME. Mime types determine, among other things, what application will handle particular file types. You can read more about mime types in Section 17.4, *GNOME Mime Types*

  - **Run Command** - This menu item allows you to run a command from GMC.

  - **Run Command in panel** - This menu item lets you run preloaded commands or commands you specify within the directory that you are currently in. Such commands might be to find SUID or SGID programs, etc.

  - **Exit** - This will allow you to exit the GNOME File Manager.

IMPORTANT

This will exit all GNOME File Manager processes which include the GNOME Desktop. If you exit you will lose all functionality on your desktop. This action is not recommended. If you wish to close the current window, use the **File/Close** button.

14.10 A Drag and Drop Tour of GNOME

There are many tips and tricks to the Desktop in GNOME. The Drag and Drop functionality extends to many areas of GNOME making it easy to interconnect GNOME in many interesting ways. Below is a series of tips and tricks to using GNOME Drag and Drop. This is a good tour of GNOME and will show you how to utilize GNOME to its fullest extent.

- Drag a Color onto the Panel - Whenever you have a color selector displayed you may drag a color from the selected color bar to the Panel and it will change the Panel to that color.

- Drag a Pixmap to the Background Selector - If you would like to change the background to an image, you can drag that image from your GNOME File Manager to the Monitor Image in the Background Capplet of the Control Center and it will change to that image.

- Drag to an Application - Many GNOME compliant applications will accept drag and drop. If you would like to open a file in Gnumeric, a GNOME compliant spreadsheet application, you may simply drag the file from the GNOME File Manager onto Gnumeric and it will open the file. The same is true for applications built using Motif. You may drag a saved URL onto Netscape 4x and it will open the URL. This can be very useful if you are working within the GNOME File Manager and wish to quickly open a file.

- Adding an Application Launcher to the Panel - If you would like to add an application launcher to the Panel you may drag and drop any executable file from the GNOME File Manager, or the Desktop, onto the Panel. This will display the Create Launcher applet dialog box which will allow you to select a name and an icon for that launcher.

- Dragging Files - There are many ways to use drag and drop to help you manage your system. You can open two GNOME File Manager windows selecting two different directories then drag files between the two windows to copy, move, or link files. You can drag files from the File Manager to the desktop to make it more accessible. Use the middle mouse button or the right and left mouse buttons together and Drag a directory folder to the desktop. Choose the link option from the pop-up menu to make a link to the desktop. This will give you a quick way to launch the File Manager to that directory.

- Dragging Directories - You can drag a directory out of the GNOME File Manger and place it on the Panel. This will create a new menu which allows you easy access to the files in that directory.

- You may drag any sub menu from the Main Menu to the panel and a new menu launcher is added to the panel. This allows easier access to that sub-set of menus.

15   Configuring the Panel

15.1 Introduction

The GNOME Panel is highly configurable and comes equipped with many graphical tools to help you do the configuration. In this section you will learn how to configure any GNOME Panel the way you would like it.

15.2 Global Panel Properties

To start configuring the GNOME panel right click on the panel and select the **Global Properties** menu item. You may also press the **Main Menu** button and select the **Panel** | **Global Properties** menu item.

This will bring the Global Panel Configuration dialog up. With this dialog you can set properties that will affect all Panels you use now and any panels you add in the future.

Figure 15–1   The Global Panel Configuration Dialog

The Global Panel Configuration Dialog contains six tabs which help you configure the global properties of the GNOME Panel: Animation, Launcher Icon, Drawer Icon, Menu Icon, and Miscellaneous. Each of these tabs is explained below.

15.2.1 Animation Tab

- **Enable animations** - This allows the animations configurable on this tab to be visible. The animations must be enabled to use the other features of this tab. The default position is on.

- **Auto-Hide Animation Speed** - If you have the panel hiding automatically this will control how fast it occurs.

- **Explicit-Hide Animation Speed** - This controls the hide speed when you press the Panel's Hide Button.

- **Drawer Animation Speed** - If you use a Drawer panel this will control how fast the drawer menu will raise.

- **Auto-Hide Minimize Delay(ms)** - If you have the Panel set to minimize automatically this will allow you to control how much time passes before it minimizes. The Panel will start the time count once the mouse is no longer over it. It will appear again when the mouse is passed over the portion of the Panel that remains visible. This time is measured in milliseconds.

- **Auto-Hide Minimized Size(pixels)** - If you have the Panel hiding automatically this determines the number of pixels that will show while the Panel is minimized.

15.2.2 Launcher icon Tab

- **Tiles enabled** - This checkbox will enable background tiles for all icons on the Panel.

- **Tile filename (up)** - This is the name and path of the image file you wish to use for the tile in the up position (inactive, not pressed) You may press the **Browse** button to search for the file. Launcher tiles must be enabled to access this option.

- **Tile filename (down)** - This is the name and path of the image file you wish to use for the tile in the down position (active, pressed). You may press the **Browse** button to search for the file. Launcher tiles must be enabled to access this option.

- **Border width (tile only)** - This determines the width of the border around an icon. This is very useful if you have an icon that would normally cover up a tile. You can set the border to a smaller size and still be able to see your tile.

- **Depth (displacement when pressed)** - This determines the depth an icon will displace when pressed. Launcher tiles must be enabled to access this option.

15.2.3 Drawer icon Tab

- **Tiles enabled** - This checkbox will enable background tiles for all drawers on the Panel.

- **Tile filename (up)** - This is the name and path of the image file you wish to use for the tile in the up position(inactive, not pressed) You may press the **Browse** button to search for the file. Drawer tiles must be enabled to access this option.

- **Tile filename (down)** - This is the name and path of the image file you wish to use for the tile in the down position(active, pressed) You may press the **Browse** button to search for the file. Drawer tiles must be enabled to access this option.

- **Border width (tile only)** - This determines the width of the border around a tile. Drawer tiles must be enabled to access this option.

- **Depth (displacement when pressed)** - This determines the depth a tile will displace when pressed. Drawer tiles must be enabled to access this option.

15.2.4 Menu icon Tab

- **Tiles enabled** - This checkbox will enable background tiles for the **Main Menu** button on the Panel.

- **Tile filename (up)** - This is the name and path of the image file you wish to use for the tile in the up position (inactive, not pressed) You may press the **Browse** button to search for the file. Menu tiles must be enabled to access this option.

- **Tile filename (down)** - This is the name and path of the image file you wish to use for the tile in the down position(active, pressed) You may press the **Browse** button to search for the file. Menu tiles must be enabled to access this option.

- **Border width (tile only)** - This determines the width of the border around a tile. Menu tiles must be enabled to access this option.

- **Depth (displacement when pressed)** - This determines the depth a tile will displace when pressed. Menu tiles must be enabled to access this option.

15.2.5 Logout Icon Tab

- **Tiles enabled** - This checkbox will enable background tiles for all Logout buttons on the Panel.

- **Tile filename (up)** - This is the name and path of the image file you wish to use for the tile in the up position (inactive, not pressed) You may press the **Browse** button to search for the file. Logout button tiles must be enabled to access this option.

- **Tile filename (down)** - This is the name and path of the image file you wish to use for the tile in the down position (active, pressed) You may press the **Browse** button to search for the file. Logout button tiles must be enabled to access this option.

- **Border width (tile only)** - This determines the width of the border around a tile. Logout button tiles must be enabled to access this option.

- **Depth (displacement when pressed)** - This determines the depth a tile will displace when pressed. Logout button tiles must be enabled to access this option.

15.2.6 Miscellaneous Tab

- **Tooltips enabled** - This enables tooltips for items on the panel. Tooltips are pop-up help dialogs that appear when you mouse is over an element on the panel.

- **Show Small Icons** - This will show small icons in the Main Menu.

- **Show ... buttons** - This will enable three small dots to appear on Main Menu items that launch dialogs.

- **Show popup menus outside of panels** - When this button is on it allows popup menus to appear away from the Panel. When toggled off, the popups will appear over the Panel. This can be useful on smaller screens or cluttered desktops.

- **Keep menus in memory** - This will keep your menus in memory so that they do not rescan for added items. This can increase the speed of GNOME, but may also result in you missing new items added to your menu.

- **Switched movement** - This allows launcher buttons on the panel to switch places with other icons when being moved.

- **Free Movement (doesn't disturb other applets)** - This feature locks the arrangement of your icons on the Panel (except the on you are moving). This is a good feature to use if you like the way your icons are arranged.

- **Prompt before logout** - This will bring up a yes/no dialog that asks you if you would really like to log out.

- **Raise panels on mouse-over** - If you are using a window manager that is not GNOME compliant it will not understand its relationship with the Panel. This can cause your Panel to be covered by applications. If you enable this feature you can have the Panel automatically raise when your mouse is over it.

- **Keep panel below windows** - If you are using a GNOME compliant window manager the window manager will understand its relationship with the Panel. If you choose this feature the window manager and GNOME will allow applications to appear over the Panel. This can be useful on smaller screens.

- **Close drawer if a launcher inside it is pressed** - By default drawers will remain open when you select an item within one. This can be annoying as the drawer will remain open until you close it with a mouse click. With this option selected drawers will close automatically when you select any item within one.

- **Applet padding** - This changes the amount of space (padding) between icons and applets.
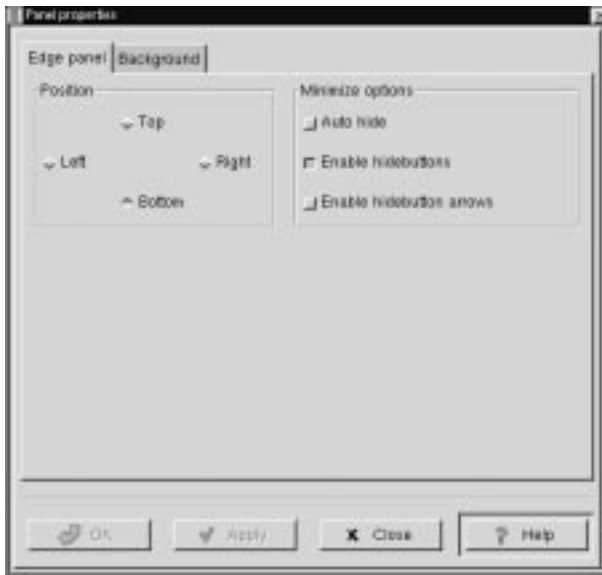
15.3 This Panel Properties

Each Panel's properties can be configured individually. To change the configuration of the active Panel, right mouse click on the Panel and select the **This panel properties**

menu item from the **pop-up** menu. You may also press the **Main Menu** button and select the **Panel** | **This panel properties** menu item.

This will bring up the Panel properties dialog box. In this box you can change the properties for the active Panel.

Figure 15–2   This Panel Properties Dialog

The Panel properties dialog contains two tabs to help you set the active Panel properties: Edge Panel and Background. Both of these tabs are explained below. If you have converted the panel to a corner panel, which is accomplished by selecting this item in the main menu, a Corner Panel tab will replace the Edge Panel tab.

15.3.1 Edge Panel Tab

- **Position** - This changes the position of the Panel on the screen. You may select either Top, Right, Left or Bottom. The Panel will change position once you have pressed the **Apply** button.

- **Minimize Options** - The options here will enable you to either Explicitly hide the Panel yourself using the hide buttons or have it Auto Hide when the mouse is not over the Panel. If you choose to Auto Hide you might want to Disable the hide buttons here as well. You may also disable the hide button arrows graphics on the Panel as well.

15.3.2 Background Tab

- **Background** - These options allow you to change the background of the Panel itself. You may choose, if you wish, to have the standard, Pixmap, or Color background. The standard look for the Panel is determined by the GTK theme you are running at the time. The Pixmap option allows you to choose an image to tile or scale to the Panel. The Color option allows you to specify a particular color for the Panel.

IMPORTANT

An easier way to change the background of your panel is to drag and drop an image file from the GNOME File Manager onto the Panel. This will automatically change the background of the panel to that image.

- **Image file** - If you wish to have a Pixmap for the background of your Panel this section of the dialog is where you choose which image to use. If you press the **Browse** button you can search for the file you want to use.

- **Scale image to fit panel** - This allows a background image to scale to the size of the panel. If it is not checked, images will tile to the panel.

- **Background color** - If you choose to have your Panel one color this button will launch a dialog which allows you to specify which color to use.

16 Editing the Main Menu

16.1 Introduction

The Main Menu is a repository for your applications and can be found on the Panel. The Main Menu is preloaded with GNOME but is can be configured to fit your work habits. The Main Menu is broken up into two main subdirectories: the System menus and the User menus. The Menu Editor is available for you to add new applications to the Main Menu but you can not add applications to the System menus if you are not system administrator (root). In this section you will learn how to configure the Main Menu with applications you wish to use everyday.

16.1.1 Configuring the Main Menu

If you want to change properties of the Main Menu or any other menu you have on your Panel you may right click on the menu button and select **Properties**. This will launch the **Menu properties** dialog.

Figure 16–1   Menu Properties Dialog
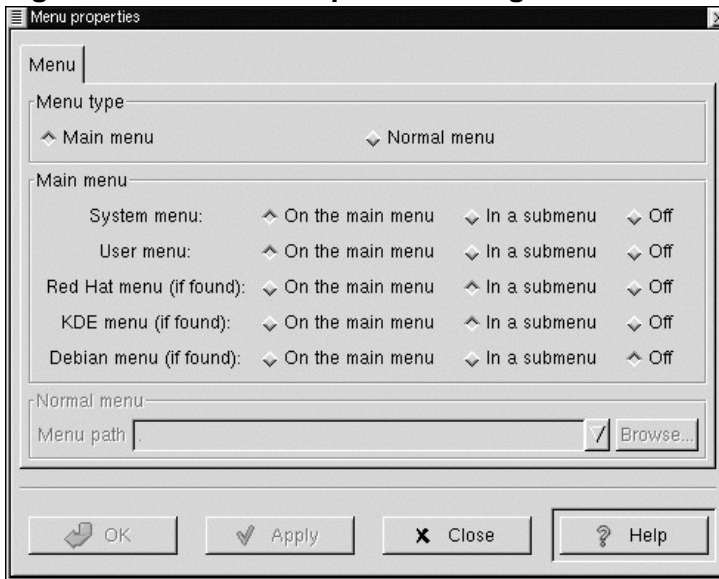
The first selection item in the **Menu properties** dialog is the **Menu type**. This will allow you to choose for each item below whether it should be directly in the Main Menu, moved to a submenu of the Main Menu, or removed from the Main Menu entirely.

The other choices in the **Menu properties** dialog are sub-menu selections for Main Menus. These selections allow you to choose what is in your Main Menu which you can have in the Main Menu, in a submenu, or off:

- System Menus - These are the menu items that are the default applications that come with GNOME.

- User Menu - This contains any menu items you added for your user account using the Menu Editor.

- Another Level Menu - If you are using the Red Hat Linux distribution this is the set of applications that ship with it by default.

- KDE Menu - If you are a user of the K Desktop Environment, you may choose to include the applications included in the KDE menus.

- Debian Menu - If you are using the Debian Linux distribution, this is the set of applications that ship with it by default.

16.1.2 Using the Menu Editor

The Menu Editor is a configuration tool for the Main Menu. It is very useful in setting up your system to your requirements. The Menu Editor is started by clicking on the Main Menu Launcher and selecting **Settings** | **Menu editor** from the **Main Menu**.

This will launch the Menu Editor.

Figure 16–2   The Menu Editor
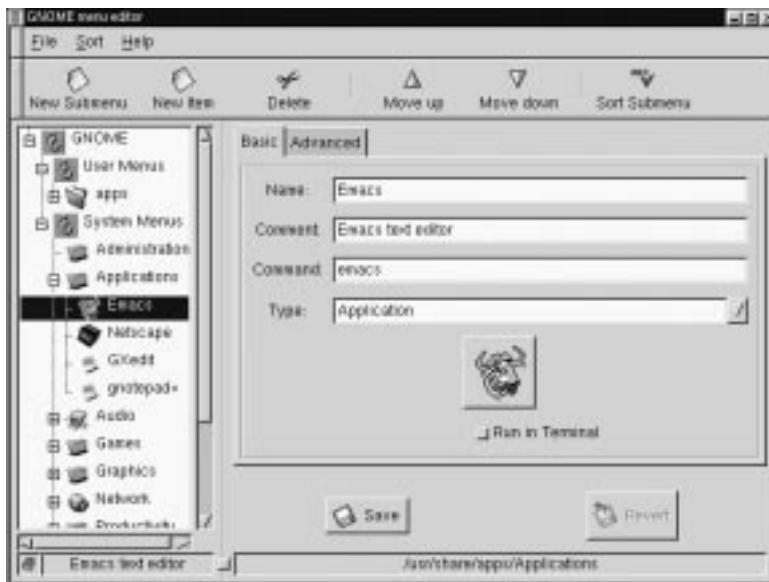
The Menu Editor is divided into two main panels. The left side contains the menu in its default state. The right side contains a tabbed dialog that allows you to add new applications to the menu.

On the left side in the menu tree you will notice that there are two main menu lists, one for User Menus and one for System Menus. The User Menus are for the current user

and the System Menus are for all users on the system. The prepackaged applications are all located in the System Menus.

Within the menu list on the left side you may open and close folders and see what is in your current menu by clicking on the small plus signs beside the menus.

16.1.3 Adding a New Menu Item

If you want to add a new menu item press the **New Item** button on the toolbar. A new menu item will be added to the highlighted menu is. If you do not have a menu highlighted, it will be placed at the top of the menu tree. Select the new item and type in the information for the item in the dialog on the right. Once the information is complete press the **Save** button and the new menu item will be inserted at the highlighted locationon the left side menu tree. You may then move the menu item by pressing the up or down buttons on the toolbar. You may also move the menu item by dragging it with your left mouse button.

IMPORTANT

Keep in mind that you can not change the System menus unless you are logged in as root. An ordinary user can onlu add to, delete from, and edit the User Menus.

16.1.4 Drag and Drop in the Main Menu

In the Menu Editor - The Menu Editor supports drag and drop functionality, which will make your work easier. You may drag applications to the folders you wish them to reside in or re-arrange you folders completely.

To the Panel - If you would like to place a menu item onto the Panel, you can drag and drop from the menu to the Panel and it will place a launcher there with all the appropriate properties set for you. If you prefer not to use drag and drop you may also right click on the menu item and choose the **Add this launcher to panel** from the **pop-up** menu.

17 The GNOME Control Center

17.1 Introduction

The GNOME Control Center allows you to configure various parts of your system using a collection of tools called "capplets". These capplets may be associated with the core set of GNOME applications or other applications for which the developers have written capplets.

Your Control Center may contain more capplets than are documented here depending on the applications installed on your system.

The Control Center is divided into two main sections, the menu of configurable capplets and the main work space.

Working with the Control Center simply requires you to select a capplet from the menu on the left and double click on it. Once this is done, the workspace will change, allowing you to configure the item.

Figure 17–1   The GNOME Control Center
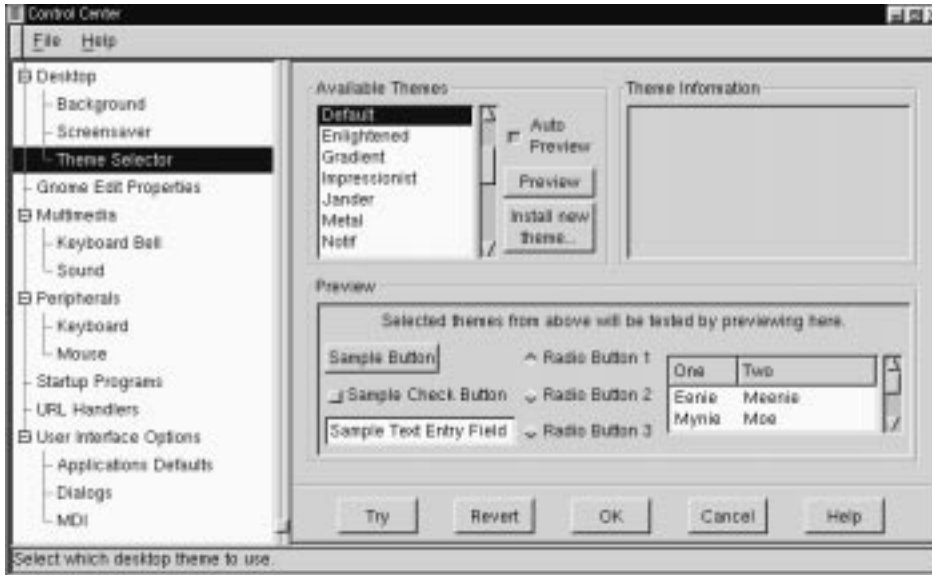
You may start the Control Center one of two ways. To launch the Control Center without any particular active capplet select the **Control Center** menu item from the **Settings** submenu in the Main Menu.

If you know which capplet you would like to edit you may start that capplet by selecting the correct menu item in the **Control Center** menu.

17.2 Desktop Capplets

17.2.1 The Background Properties Capplet

The properties for your background image can be set here by either selecting a color or an image. If you select a color you have the option of having Solid or Gradient colors. If you choose to have Gradient colors you may choose between a Vertical or Horizontal gradient and choose the second color for the gradient to end on.

If you decide to have an image as wallpaper you may browse for the image you wish to use. Once you have found your image you need to decide whether you would like to have the image tiled across the screen, centered, scaled up while keeping the same aspect, or simply scaled up to fit the screen. Once you have changed your background properties you may press the **Try** button at the bottom of the Control Center to make the change.

If you would like to set the background by any other means you may disable this capplet by selecting **Disable background selection**.

Figure 17–2   The Background Properties Capplet
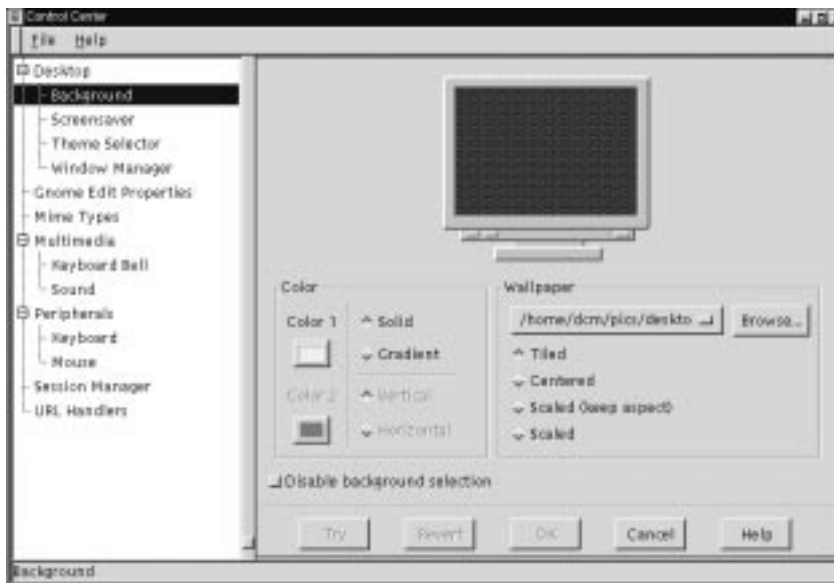
17.2.2 The Screensaver Capplet

In this capplet you can change your xscreensaver properties. This capplet contains a list of available screen savers you may choose and a demo screen. Below these two dialogs you will see a set of tools that allow you to change the settings for the global

screen saver properties. If the screen saver you choose has particular settings you can change those by pressing the **Settings** button that appears below the Screen Saver list.

Global Screen Saver settings - In this section of the capplet you can change the time, password, and power management properties. You can decide how long you would like the screen saver to wait before starting by typing the number of minutes in the **Start After** text box. If you would like a password to return to your desktop click the **Require Password** button. Your account login password is the password set for the screen saver.

You are also given the option of using power management -- if your monitor is capable of it. You may set the time to wait before the monitor is shut off by typing the time in the **Shutdown monitor** text box.

Figure 17–3   The Screensaver Properties Capplet
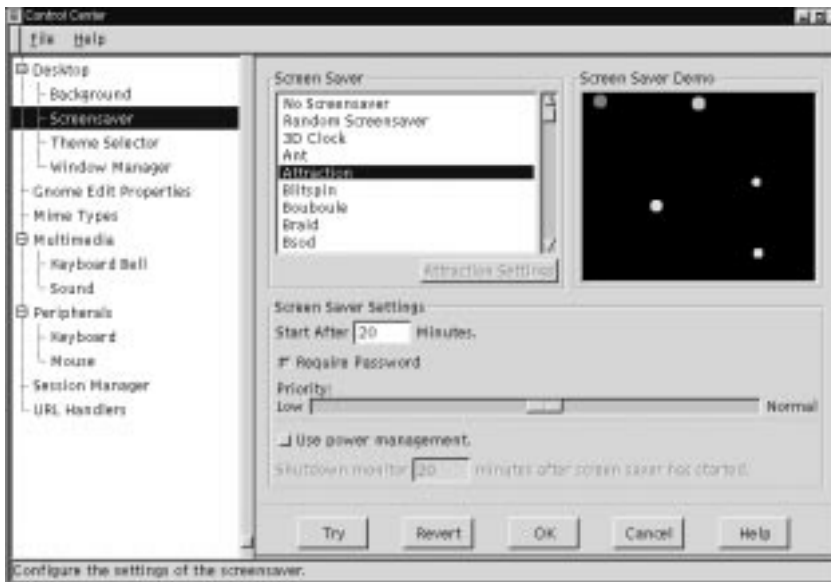
17.2.3 Theme Selector

The Appearances capplets contains the Theme Selector which allows you to select which GTK theme you would like to run.

GTK themes are coordinated settings that allow the GTK widget set to change its look and feel. The widget set is the set of tools that provide buttons, scrollbars, checkboxes, etc. to applications. GNOME compliant applications use the GTK tool set so most of your GNOME applications will change their look and feel if you change the GTK theme.

To change your GTK theme select a theme from the Available Themes list on the left side of the main workspace. If you have Auto Preview selected you will be able to see what the theme looks like in the preview window below. If you like the theme press the **Try** button on the bottom of the GNOME Control Center to install it.

There are a few GTK themes that come loaded with GNOME when you install it. If you would like more themes you can check resources on the Internet like http://gtk.themes.org. Once you have found and downloaded a theme you like, press the **Install new theme** button. This will launch a file browser that allows you to find the theme you have just downloaded. The theme files should be in a `tar.gz` or `.tgz` format(otherwise known as a "tarball"). Once you have found the file press the **OK** and gnome will install the theme for you automatically. Now you can look in the Available Themes list for the theme you have installed.

Once the theme has been unpacked into the .themes directory it will be listed in the available themes window the next time you start the GNOME Control Center.

Figure 17–4   The Theme Selector Capplet

If you would like to change the font used in the current theme you may do so by selecting the **Use custom font** checkbox and selecting the font from the font button below it. This will bring up a font selection dialog that allows you to specify the font, its style, and size.

17.2.4 Window Manager Capplet

Because GNOME is not dependent on any one window manager this capplet allows you to select which window manager you wish to use. The Window Manager capplet does not determine which window managers you have available but allows you to define what, and where they are.

Figure 17–5   Window Manager Capplet
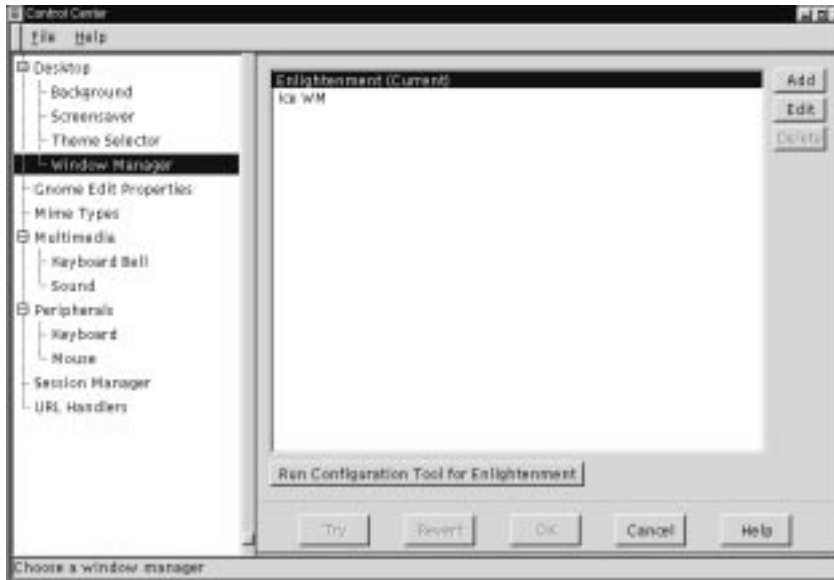
The Window Manager capplet has a main list of the window managers that you can currently select from. Whichever window manager is active is notated by the word "Current".

If you wish to add a new window manager to the main list you may press the **Add** button. This will launch the **Add New Window Manager** dialog.

Figure 17–6   Add New Window Manager

In the **Add New Window Manager** dialog you may specify the name you wish to give the window manager, the command to launch that window manager, and the command to launch any configuration tool that might be available for that window manager.

If you know that the window manager is fully GNOME compliant and can be session managed you may select the **Window manager is session managed** button. If you are unsure you should check the documentation of your window manager.

Press **OK** when you are done.

Once you have finished adding your new window manager you will see it appear in the main list of window managers. If you need to change any of the properties you set in the **Add New Window Manager** dialog you may select the window manager from the main list with your mouse and press the **Edit** button.

You may also delete any window manager in the main list by selecting it with your mouse and pressing the **Delete** button.

If you are ready to switch the current window manager you may select the manager you wish to run from the main list and press the **Try** button. If you would like to run the configuration tool, make sure the manager you want to configure is selected and press the **Run Configuration Tool for [window manager name]** button.

17.3 Default Editor

The Default Editor Capplet allows you to select which editor will be your default editor while using GNOME. This will allow applications like the GNOME File Manager to launch the correct editor when you try to open files associated with editing. All popular editors available are included in the selection list. This Capplet is very similar to the Mime Type Capplet but is used in association with certain applications.

Figure 17–7   The Default Editor Capplet
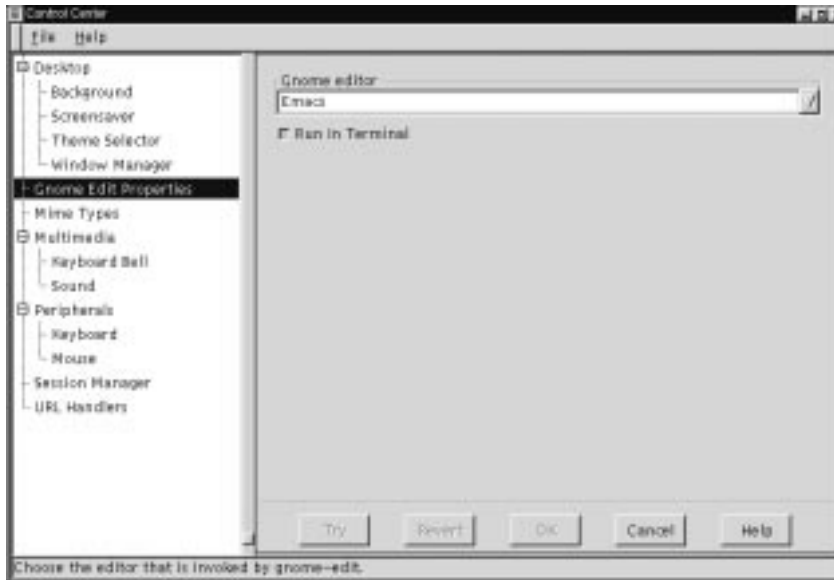
17.4 GNOME Mime Types

The GNOME Mime Types Capplet allows you to determine how you wish to handle certain file types, or Mime types. Mime stands for Multipurpose Internet Mail Extensions and was originally developed to allow email to carry various forms of data. In GNOME you can define certain Mime types to be handled in any manner you wish. For example, if you use .sgml files frequently and you wish to always use Emacs to edit them you can configure the .sgml Mime type to always be handled by Emacs. This means that any program that wishes to launch the mime type for you will bring up Emacs. This includes double clicking on the file type in the GNOME File Manager.

Figure 17–8   The GNOME Mime Types
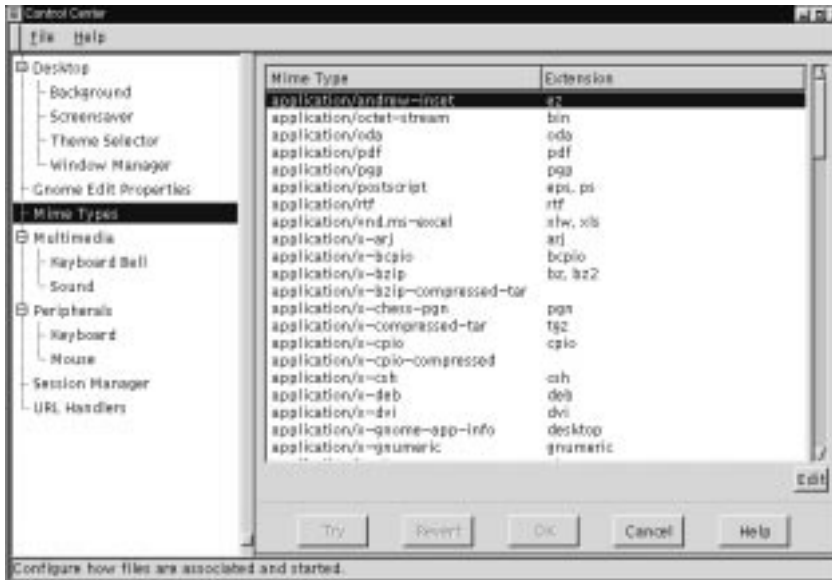
To add a new mime type press the **Add** button. This will display the **Add New Mime Type** dialog in which you may define the new Mime type.

To edit an existing Mime type you may select the Mime type with your mouse cursor and press the Edit button. This will bring up the **Set Actions for...** dialog. You may define the icon used for the Mime type, the **Open** action, the **View** action and the **Edit** action.

17.5 Multimedia Capplets

17.5.1 Keyboard Bell

The Keyboard Bell capplet allows you to change the bell sound that is produced by your CPU speaker when a keyboard error or message is sent.

Volume changes the actual volume of the bell.

The pitch slider will change the pitch of the note that is played. By default it is set to 440Hz, or the A above middle C.

Duration changes the length of time the tone is played.

The **Test** button will allow you to hear the current settings of your keyboard bell.

Figure 17–9   The Keyboard Bell Capplet
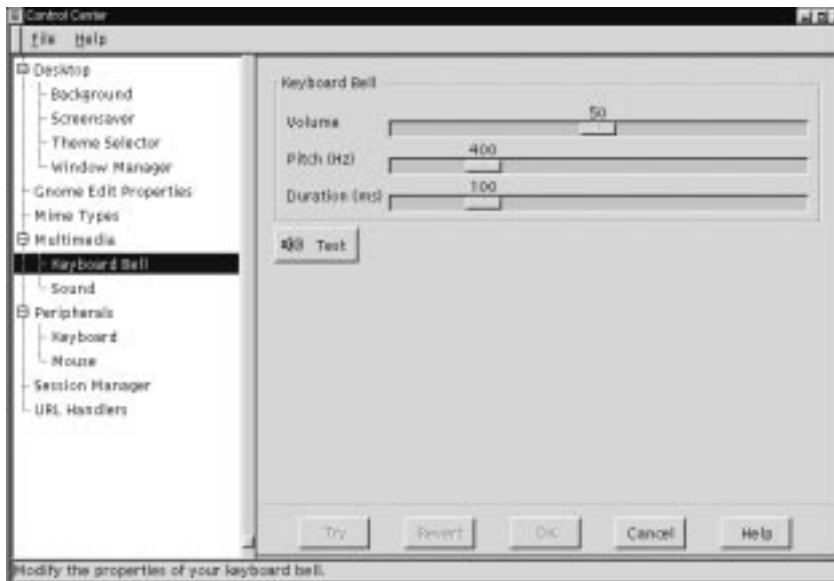
17.5.2 The Sound Capplet

The Sound capplet allows you to set the system sounds for your GNOME session. There are two tabs to select in the sound capplet: General and Sound Events.

General Tab - At this point you have two options to choose from in the General tab, enabling sound for GNOME and for events. If you select **Enable Sound for GNOME** you will make sure that GNOME's sound engine (ESD) will be launched every session of GNOME you run. **Enable sound for events** will launch any sound files you have set in the **Sound Events** tab when those events occur. With these two items enabled

you will utilize GNOME's session management, which will remember your sound settings whenever you log in or out.

Sound Events Tab - This tab allows you to navigate through the sound events in GNOME and change their sounds.

To change a sound associated with a GNOME event select the event in the hierarchical list on the left and press the **Browse** button to find a sound file on your system that you wish to associate with that event. Once you have found a sound file you may press the **Play** button to test the sound and see if you like it enough to hear it every time the event occurs.

Figure 17–10   The Sound Properties Capplet
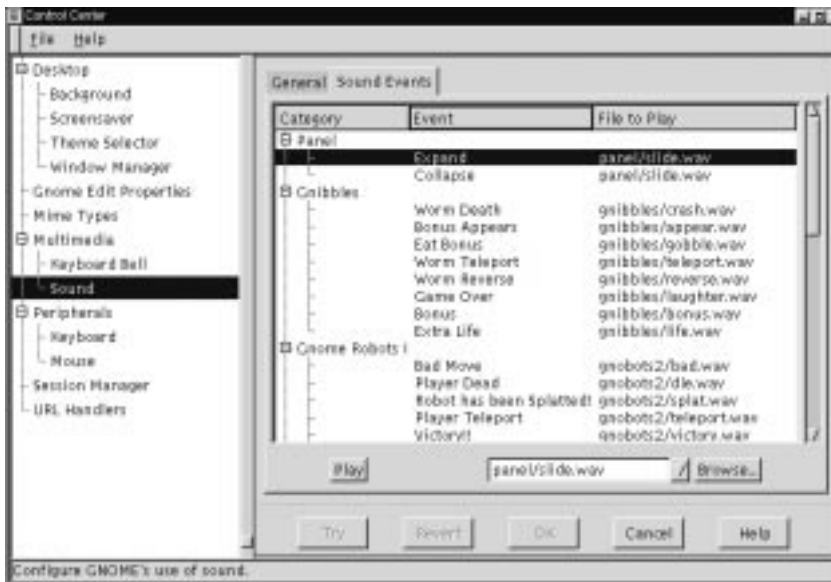
17.6 Peripherals

The capplets in this section of the Control Center will help you configure hardware input devices including keyboard, and mouse properties.

17.6.1 The Keyboard Properties Capplet

There are currently two settings for the keyboard in this capplet. You may change the properties of **Auto-repeat** and the **Keyboard Click**.

Auto-repeat enables you to hold a key down and have it repeat the character at the rate and delay you set in this capplet. **Keyboard Click** enables a small click sound to play at each key press.

Figure 17–11   The Keyboard Properties Capplet
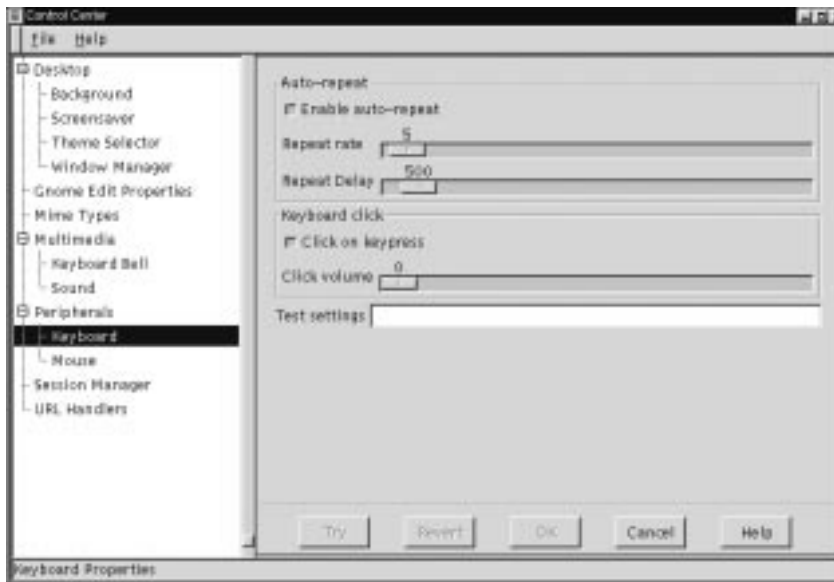
17.6.2 The Mouse Properties Capplet

The Mouse Properties capplet allows you to change between left and right handed mouse buttons and to define the **Accelerations** and **Threshold** properties.

The **Accelerations** setting allows you to change the speed the mouse moves across the screen in relation to the movement of the mouse on your mouse pad. The **Threshold**

setting allows you to set the speed at which you have to move your mouse before it starts the acceleration speed you have defined in the **Acceleration** setting.

Figure 17–12   The Mouse Properties Capplet

17.7 Session Manager

The Session Manager Capplet allows you to control the GNOME Session Management. This includes which programs start up when you log in, whether to automatically save your GNOME requests, and whether to confirm log out requests. You can find out more information about Session Management in Chapter 18, *A Word About Session Management*.

- **Prompt on logout** - This first option allows you to disable the confirmation dialog when logging out.

- **Automatically save changes to session** - This will make the Session Management always save changes made to your GNOME session when you log out. If this option is not set you can still have GNOME save your session changes by checking the appropriate box in the log out confirmation dialog.

- **Non-session-managed Startup Programs** - This allows you to start non-session managed applications whenever you start a GNOME session.

NOTE

Programs that are not GNOME compliant are not session managed so will not be restarted without being listed here. You do not need to put GNOME applications in here, you can simply leave them running and save the current session when you log out.

If you wish to add a new program to the **Non-session-managed Startup Programs** press the **Add** button. This will launch a simple dialog that allows you to specify the command to launch the application and what priority it will receive.

The priority for most applications you wish to start is 50. If you have an application that needs to be started before other applications, like a window manager, you should set the priority to a lower number.

IMPORTANT

This option is for advanced users. Unless you are familiar with the Priority settings you should keep you applications running with a Priority of 50.

- **Browse Currently Running Programs** - This allows you to see what applications are currently running. You can shut down those applications if you wish and those applications will be removed from your GNOME session. The applications in this list are mostly higher level applications and should not be shut down. However, if there is a part of GNOME that you do not wish to run, like the Panel, this is

where you would shut it down for now and prevent it from being started when you initiate GNOME in the future.

IMPORTANT

This option is for advanced users only. You should not use this tool shut down applications you may wish to use the next time you log in to GNOME.

Figure 17–13   The Session Manager Capplet
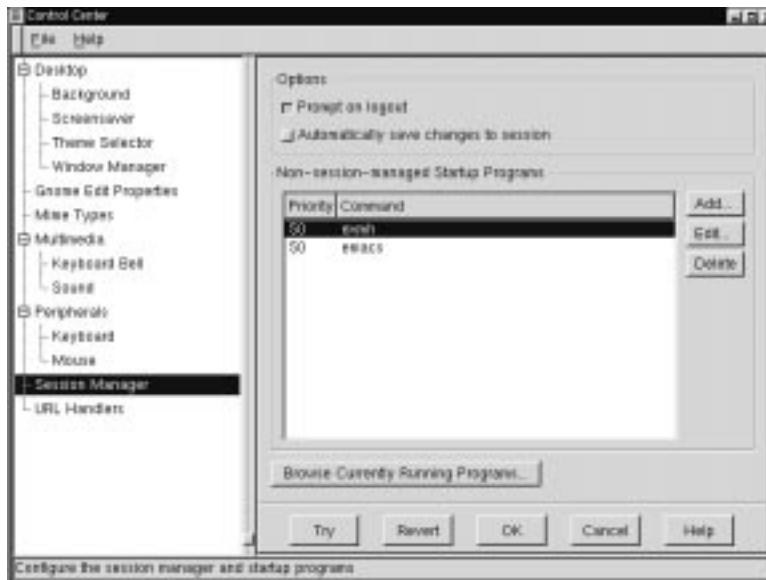
17.8 User Interface Options

The User Interface Options allows you to change the appearance of applications that are GNOME-compliant. You may recognize these applications as ones that are pre-installed with GNOME or ones that say they are built with GTK(the GIMP Toolkit).

17.8.1 Application Defaults

The Application Defaults capplet allow you to change certain user interface aspects of your GNOME compliant applications.

IMPORTANT

Although this capplet gives you great control over the look and feel of your applications you should consider these tools for advanced use only.

- Can detach and move menubars - By default menubars in GNOME applications may be dragged from their usual location and placed anywhere within the application or desktop. If you do not wish to use this feature you may turn it off.

- Menus have relieved borders - By default menubars have relieved borders. If you do not like this look you may turn this feature off.

- Submenus can be torn off - This allows the submenus to have the perforated line which allows you to "tear" them off an have them as a small movable window.

- Menu items have icons - Some menu items in applications will have icons. If you wish not to see these icons in applications that use them you may turn off this feature.

- Statusbar in interactive when possible - Some applications can have the status bar at the bottom become separated into its own window. If you would like to have those applications separate the status bar into another window you may turn on this option.

- Statusbar progress meter on right - Some applications have progress meters in their statusbars. By default these progress meters are on the right side of the statusbar. If you wish them to be on the left you may turn off this feature.

- Can detach and move toolbars - By default toolbars in GNOME applications may be dragged from their usual location and placed anywhere within the application or desktop. If you do not wish to use this feature, you may turn it off.

- Toolbars have relieved border - By default toolbars have relieved borders. If you do not like this look you may turn this feature off.

- Toolbar buttons have relieved border - By default toolbar buttons do not have relieved borders in their natural state. They do, however, change when the mouse is over them. If you wish them to be relieved at all times you may turn on this feature.

- Toolbars have line separators - By default toolbar buttons have small line separators between them. If you so not wish to have the line separators you may turn this feature off.

- Toolbars have text labels - By default toolbar buttons have images and text to identify them. If you are familiar with the buttons and do not need the text you may turn on this feature.

Figure 17–14   Applications Defaults Capplet
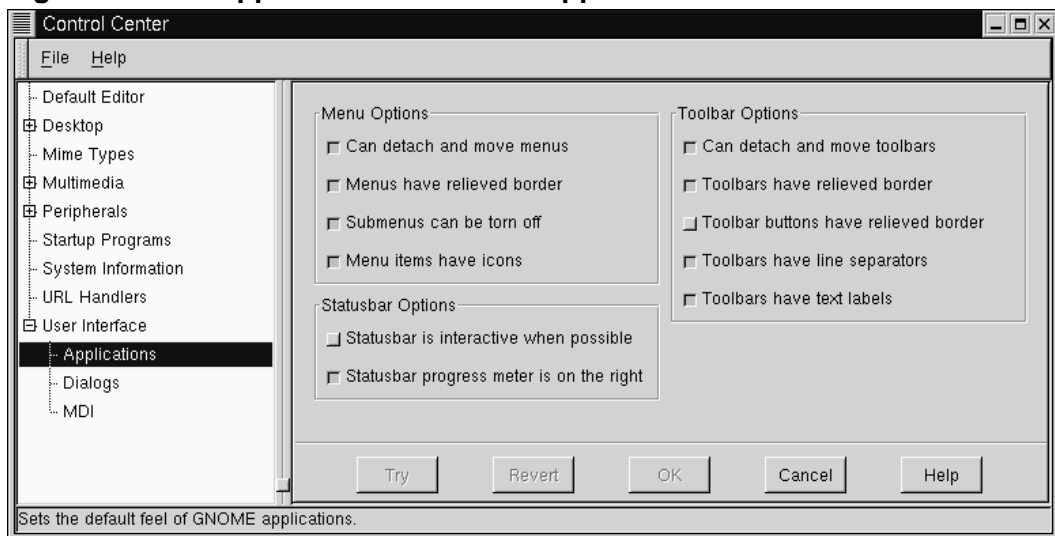
17.8.2 Dialogs

The Dialogs Capplet will allow you to change the default settings for dialog boxes in GNOME compliant applications. A dialog box is a window that is launched by

an application to help perform a task needed by that application. An example of a dialog box is a Print dialog which appears when you press a print button. The dialog allows you to set print options and start the print process. The Dialogs capplet will allow you to change the following options:

Dialog buttons - Choose to use the default buttons, buttons more spread out, put buttons on the edges, put the buttons on the left with left-justify, and put buttons on the right with right-justify.

Dialog buttons have icons - Some dialog buttons (for example "OK") can have icons on them. By default the applications which provide this have the icons turned on. If you wish not to see them you may turn off this feature.

Use statusbar instead of dialog when possible - You may tell applications to use the statusbar instead of a dialog if the application will allow it. This will only work with dialogs that provide information not one that require some interaction on your part.

Dialog position - This will let you choose how the dialogs appear when launched. You can let the window manager decide for you (or how you have defined it in the window manager configuration), center the dialogs on the screen, or drop them where the mouse pointer is when they are launched.

Dialog hints - This will let you change the behavior of the dialog hints which are the tooltips that appear when you move your mouse button over a button or part of the dialog. You may choose to have hints handled like other windows, or let the window manager decide how to display them.

Place dialogs over application window when possible - You may choose to place dialog over the applications when possible which will help you keep your windows organized on your screen If you are familiar with other operating systems you may wish to keep this selected as this is how most operating systems handle dialogs.

IMPORTANT

Although this capplet gives you great control over the look and feel of your applications you should consider these tools for advanced use only.


Figure 17–15   Dialog Capplet
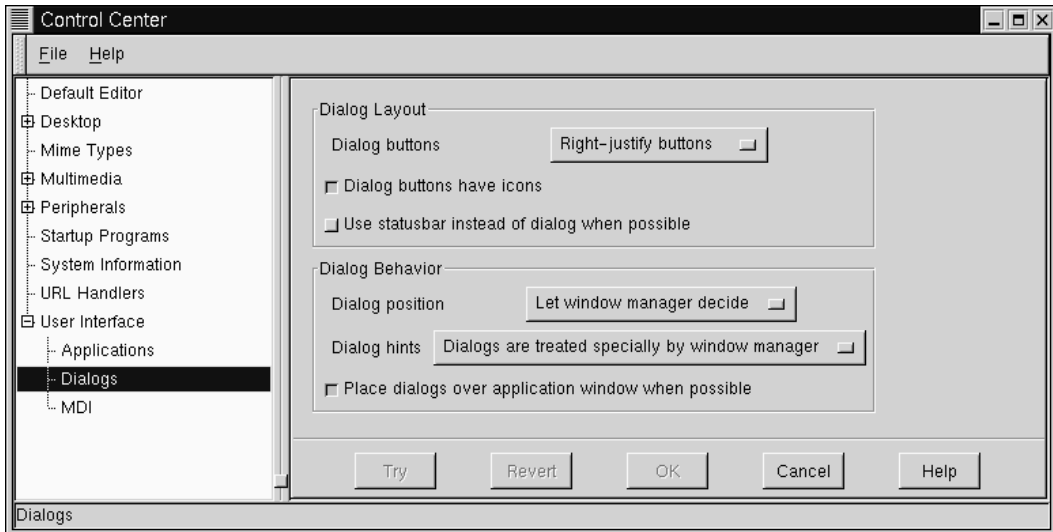
17.8.3 MDI

The MDI capplet allows you to change the MDI mode for GNOME applications. MDI stands for Multiple Document Interface and refers to the how more than one document is displayed in GNOME applications.

IMPORTANT

Although this capplet gives you great control over the look and feel of your applications you should consider these tools for advanced use only.

The default style in GNOME-compliant applications for MDI is usually tabs or "notebooks" If you do not like the tab look you may change it here.

Deafault MDI Mode - In addition to Notebook, you will find, Toplevel and Modal. Notebook is the default tab look, Toplevel displays only the active document on the

top view until it is closed and Modal has only one toplevel which can contain any of the documents at any one time, however only one can be displayed. If you have ever used Emacs Modal is very similar to the way Emacs handles buffers.

MDI notebook tab position - If you choose to use the Notebook style you may then decide where you want the tabs to appear in your applications. You may have them at the top, left, right, or bottom of your application. Keep in mind that these choices will affect applications that are GNOME compliant.

Figure 17–16   MDI Capplet
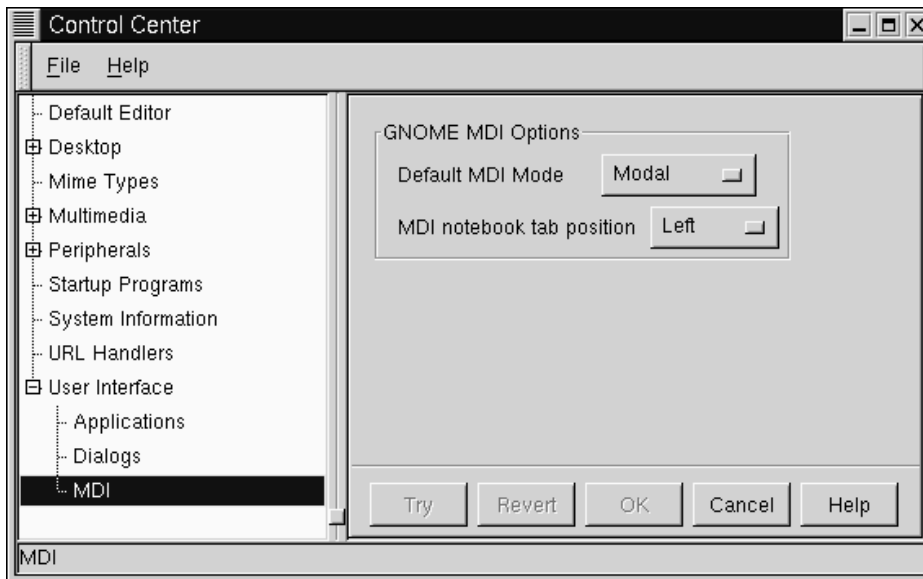
18 A Word About Session Management

18.1 Introduction

You might have seen a little bit about session management when you read about GNOME. Session management can be a very useful tool for you in your use of GNOME and GNOME applications.

The main idea of session management is that your work will be saved whenever you log out of GNOME. Your GNOME "Session" is currently saved when you log out, however you may not notice all that it can do.

Probably one of the most useful features of session management is the ability to start applications you had open when you logged out of GNOME. This is easily demonstrated as you exit GNOME and enter again. You will see your applications re-appear in the same location on your desktop as they were when you logged out.

If there are certain applications you wish to start up whenever you log in, even if they were not open the last time your session was saved, you may add them to the Session Manager Capplet which you can read about in Section 17.7, *Session Manager*.

<div align="center">

IMPORTANT

</div>

If you have heard about application data being session managed you have heard about the future of GNOME. For application data to be saved correctly you must be using a fully compliant GNOME application. At the time of publication of this manual there were not any applications that offer full session management for your data, so until there are some and you are aware of their capabilities you should not rely on session management to save your application data.

18.2 Resetting the GNOME Session

One advanced feature of the GNOME Session Manager is the ability to recover a "clean session" if anything goes wrong for you. To do this you must hold down the [CTRL] and [SHIFT] keys together when you log in to GNOME.

This will bring up a dialog box which gives you two different options for restoring your GNOME Session.

The first option is to **Start with default programs**. This option will remove all of the session configuration setting in respect to applications. This will only erase the GNOME session data for applications you had running when you logged out last, it will not change any information you may have set in the Session Manager Capplet in the GNOME Control Center.

The second option is to **Reset all user settings**. This will reset all GNOME application, and core configuration data. This option will destroy any configurations you have made to the Panel, the GNOME File Manager, the Session Manager Capplet, and any GNOME application. This option will not remove files on your desktop.

IMPORTANT

These options are for advanced users and should only be used in case there is a problem with your GNOME Session. You can lose data for many applications if you utilize the functionality provided by these options.

Figure 18–1   Login Screen for Resetting GNOME Session

19   Panel Applets

19.1 Introduction

This section describes the GNOME applets that are available to add to the GNOME Panel. To access these applets right mouse click on the Panel and select **Add new applet** from the **pop-up** menu.

19.2 Amusements

The Amusement applets are designed to make your life more enriched by providing you with a complete waste of time and resources purely for your entertainment. If you find that valuable space in your brain has been filled with something an Amusement applet has taught you, then the authors of that applet have succeeded in their task.

19.3 Monitors

Monitor Applets are designed to be used to keep track of your system and its functions. You can monitor the resources of your machine, which allows you to keep close tabs on how things are working.

19.3.1 Battery Monitor

The Battery Monitor is a simple applet that allows you to see how much time you have remaining on your laptop monitor.

19.3.2 Battery Charge Monitor

The Battery Charge Monitor indicates what percentage or your laptop's battery charge remains and how long the laptop can run with this charge. It also indicates if the laptop is currently plugged in. If you left click on the applet, the display will change to a graph plotting the remaining charge on the vertical axis and the time on the horizontal. So if your laptop has been charging, the graph should be rising and if it has been running on batteries, the graph should be falling.

Right clicking on the applet will bring up a pop-up menu. If you click on the Prop-
erties menu item you can edit properties such as which colors to use in the display,
when to notify you of a low battery, and the general look of the applet.

19.3.3 CPU/MEM Usage Monitor

The CPU/MEM Usage Monitor is an applet that will show you the current usage of
CPU, Memory, and Swap Space. The applet consists of three bars that are shaded
with colors to represent the usage. If you are running the applet on a horizontal panel
the top bar is CPU, the middle bar is Memory, and the bottom bar is Swap Space. If
you are running the applet on a vertical panel the CPU is the left bar, Memory is the
middle bar, and Swap Space is the right bar.

Figure 19–1   The CPU/MEM Applet

- CPU - The CPU bar will break down current CPU usage by means of three colors.
  The yellow shows the current hit by the current user. Grey shows the current hit
  by non-user specified system activity. Black show idle use.

- MEMORY - The Memory bar shows the current physical memory usage with
  one of four colors. The yellow show the current shared memory usage. The
  grey-yellow show other memory usage. The grey shows the buffers being used.
  The Green show free physical memory available on your system.

- SWAP SPACE - The Swap Space bar shows how much swap space is being used
  with a yellow bar. Any free swap space will be shown with a green bar.

19.3.4 CPULoad Applet

The CPULoad Applet is a simple graph that shows you the current CPU Usage using
three colors:

Figure 19–2   The CPULoad Applet
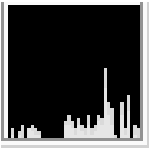
The yellow shows the current hit by the current user. Grey shows the current hit by non-user specified system activity. Black show idle use.

19.3.5 MEMLoad Applet

The MEMLoad is a simple graph that shows you the current Memory usage using four colors:

Figure 19–3   The MEMLoad Applet
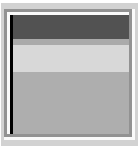
The yellow show the current shared memory usage. The grey-yellow show other memory usage. The grey shows the buffers being used. The Green shows free physical memory available on your system.

19.3.6 SWAPLoad Applet

The SWAPLoad is a simple graph that shows you the current Swap Space usage using two colors:

Figure 19–4   The SWAPLoad Applet
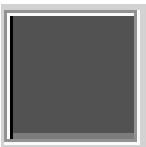
The yellow bar shows how much swap space is currently being used. Any free swap space will be shown with a green bar.

19.4 Multimedia

The Multimedia Applets are a collection of applets that allow you to utilize multimedia on your system. You will find applets that control sound, video, and other multimedia.

19.4.1 CD Player Applet

The CD Player Applet is a simple CD player that resides in the panel.

When you have an audio CD in your CD-ROM drive you can use the Play/Pause, Stop, Forward, backward and Eject buttons to control your CD for playback. Above the buttons is a small display that shows the time remaining on the track and in between the Forward and Backward buttons is the track number.

Figure 19–5   The CD Player Applet
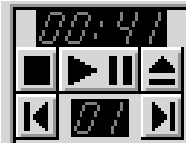
IMPORTANT

You must have the correct access rights to your CD-ROM drive for this applet to be successful. Some systems, will normally grant you the necessary rights automatically when you log into the console. The mechanism that does this is called the pam_console. If your system doesn't give you the necessary rights to the CD-ROM, then you will need to be given those rights. If you have the root password type the following in a terminal window.

```
$ su
$ Password: [type  in root password]
$ chmod a+r /dev/cdrom
$ exit
```

If your CD-ROM is located somewhere other than /dev/cdrom make sure you change it in the commands above.

If you want more control over your CD you can right mouse click on the CD Player Applet and select **Run gtcd** from the **pop-up** menu. This will launch the GNOME CD Player which you can read more about in Chapter 20, *GNOME CD Player*.

IMPORTANT

When the CD Player Applet is running you will not have physical control over the eject button on your CD-ROM drive. If you wish to eject the CD you must use the Eject Button on the CD Player Applet.

19.4.2 Mixer Applet

The Mixer Applet is a simple applet that allows you to control the volume on your system. There are two main controls on the applet: the Volume Slidebar and the Mute button. To raise or lower the volume use the slide bar with your mouse. To mute your system press the small mute button on the bottom of the applet.

Figure 19–6   The Mixer Applet

If you want more control over your system volume you can right mouse click on the Mixer Applet and select the **Run gmix** menu item from the **pop-up** menu.

19.5 Network

Network applets enable your work on a network or the internet to be monitored or enhanced. Network applets range from checking mail to monitoring the time you have spent on the internet.

19.5.1 MailCheck Applet

The Mailcheck applet will check mail on your system and let you know if there is new mail waiting for you. At this point it will only check your local system mail. It will not query another server for your mail. To access the properties for the Mailcheck applet right mouse click on the applet and select the **Properties** menu item from the **pop-up** menu. The Mail check properties dialog box will appear and you can change the following properties. If you would like to execute a program before each update you can enter a command in the Execute text box. This can be useful if you want to run fetchmail to retrieve your mail from another server. You can also specify how often you check for mail and select which animation style you would like displayed in the applet.

19.5.2 WebControl Applet

The WebControl applet allows you to launch your web browser with the URL you indicate in the URL text box. If you would like to launch a new window instead of using the active one you can check the Launch new window checkbox. To clear the URL textbox, press the **Clear** button. To access the properties, right mouse click on the WebControl applet and select the **Properties** menu item from the **pop-up** menu. The Web Control properties dialog box will appear and you can select whether you want to Display the URL text label and the launch new window options on your applet.

19.5.3 Modem Lights Applet

The Modem Lights Applet is a handy applet that will start you dial up connection and monitor it for you once it is connected. It derives its name by displaying two small modem lights which monitor the activity on your modem.

Setting up modem lights requires you to already have a dial up connection working on your system. You must right mouse click on the applet and select the **Properties** menu item from the **pop-up** menu. This will launch the Modem Lights Setting dialog.

Within this dialog are two tabs, the **General** tab and the **Advanced** tab. In the **General** you type the scripts used to start your dialup connection in the **Connect command** textbox. The disconnect scripts are typed in the **Disconnect command** textbox. Press Apply and then press the small single button at the edge of the applet. The scripts will be run. Once there is a connection the small modem lights will blink and the internet load will be displayed in the small black box in the center of the applet.

19.6 Utility

The Utility Applets are a set of general utilities to use in your work environment.

19.6.1 Clock Applet

The Clock applet is the one of the only applets that is loaded by default when you install GNOME for the first time. To access the clock properties, right mouse click on the clock and select the **Properties** menu item from the **pop-up** menu. The Clock Properties dialog box will appear allowing you to specify whether you would like 12 or 24 hour time to be displayed.

19.6.2 Printer Applet

The Printer Applet is represented by a small printer icon that lives on your Panel. If you drag a file to the Printer Applet it will print the file for you. To set up the Printer Applet, right mouse click on it and select the **Properties** menu item from the **pop-up** menu. This will bring up the Printer properties dialog box. In this dialog you can specify a printer name and the Print Command. For most systems the Print Command will be

```
lpr
```

Figure 19–7   The Printer Applet Properties
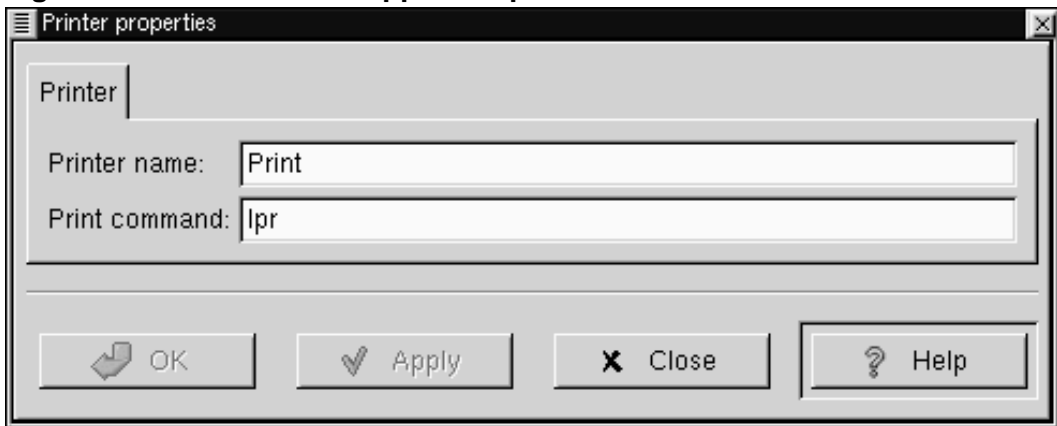
19.6.3 Drive Mount Applet

On many UNIX-like systems, after inserting a disk you must tell the computer to mount it in order to use that disk.

The Drive Mount Applet allows you or your systems administrator to mount a drive on your system by simply clicking the icon on your Panel. In order for this to work you will to have to set the drive you want to access to be user mountable.

This can be done quite easily with someone with root permissions if they have *linuxconf* installed on your machine. Just select the drive you want to access in the **Access local**

drive section. In the **Options** tab select the **User Mountable** option. This drive will now be mountable by all users.

If you do not have *linuxconf* someone with root access must edit your */etc/fstab* to include user access. This is done by adding user access to the drive specification. For Example:

If your fstab file looks like this:

```
/dev/mnt/cdrom iso9660 exec,dev,ro,noauto 0 0
```

Add "user" to the fourth column:

```
/dev/cdrom /mnt/cdrom iso9660 user,exec,dev,ro,noauto 0 0
```

Now that you can mount the drive without being root you can add the Drive Mount Applet to your panel by selecting **Drive Mount** from the **Utilities** menu in the **Add new applet** menu.

You will see a small drive image on your Panel that looks like a floppy drive.

Figure 19–8   The Drive Mount Applet

The Drive Mount applet will always default to access your floppy drive. You can change this by right mouse clicking on the applet and selecting the **Properties** item from the **pop-up** menu.

Figure 19–9   The Drive Mount Applet
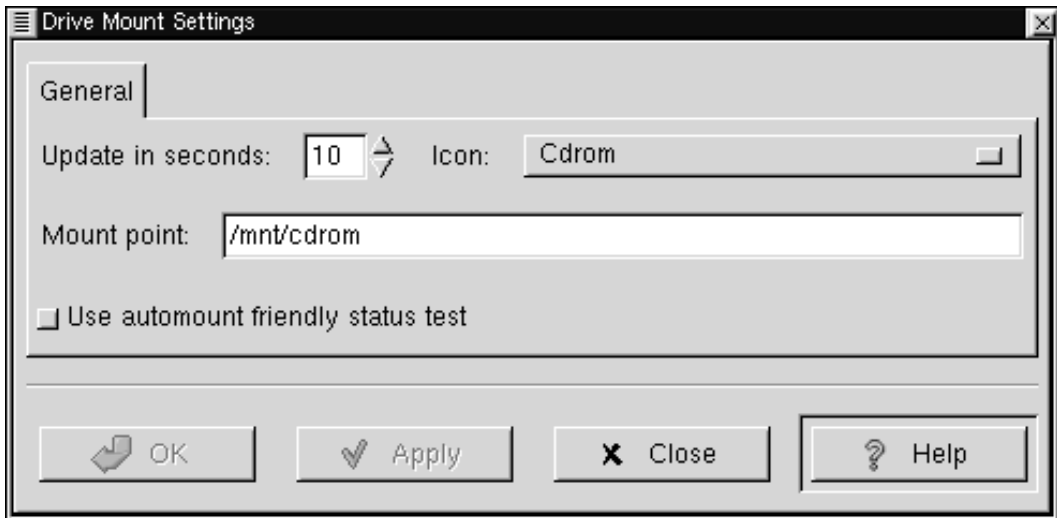
The **Drive Mount Settings** dialog allows you to define which drive you want to mount and where it is located.

The first option is how many seconds you wish to have between applet updates. During updates, the applet will check the drive status to see if it is mounted (in case it was mounted or unmounted by other means) and will display the applet accordingly.

The second option you have is which icon to be displayed. You have a choice of four icons, Floppy, CD-ROM, Zip Disk, and Hard Drive. After you have selected the icon you must put the correct mount point for the drive in the **Mount point** text box.

The last option to set in the **Drive Mount Settings** is whether or not you wish to **Use automount friendly status test** If you are using a system that utilizes autofs to auto-mount your drives the Drive Mount applet might interfere with autofs. If this is the case you should select this option. If you are not using autofs (which is most likely the case) do not select this option as it is taxing on your system and is much slower.

19.6.4 GNOME Pager

The GNOME Pager is an applet that will show you all of your virtual desktops and the applications within them. There are two main areas on the GNOME Pager, the desktops view and the applications view. In the desktop view each of your desktops will be represented as a small rectangle. If there are any applications on the desktops, they will show up as small outlines according to their position on the desktop. The Applications view will show you the applications on your active desktop in a list view. If you press the button in the middle that contains the arrow it will show you a view of all desktops and list the applications that are currently on them.

Figure 19–10   The GNOME Pager
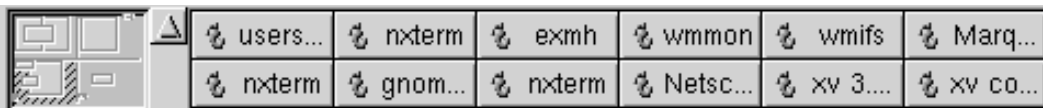
NOTE

If windows "disappear" from your screen when you iconify them, just add a GNOME Pager to your Panel.

You can access the GNOME Pager Settings dialog with a right mouse click on the Pager and select the **Properties** menu item from the **pop-up** menu.

The GNOME Pager Settings dialog contains two tabs, the **Pager** and the **Tasklist** tabs. The **Pager** tab allows you to change the setting for the Pager (graphical representation of desktops) and the **Tasklist** tab allows you to change the settings of your Tasklist (list of open applications on your desktop).

Figure 19–11   The GNOME Pager Properties
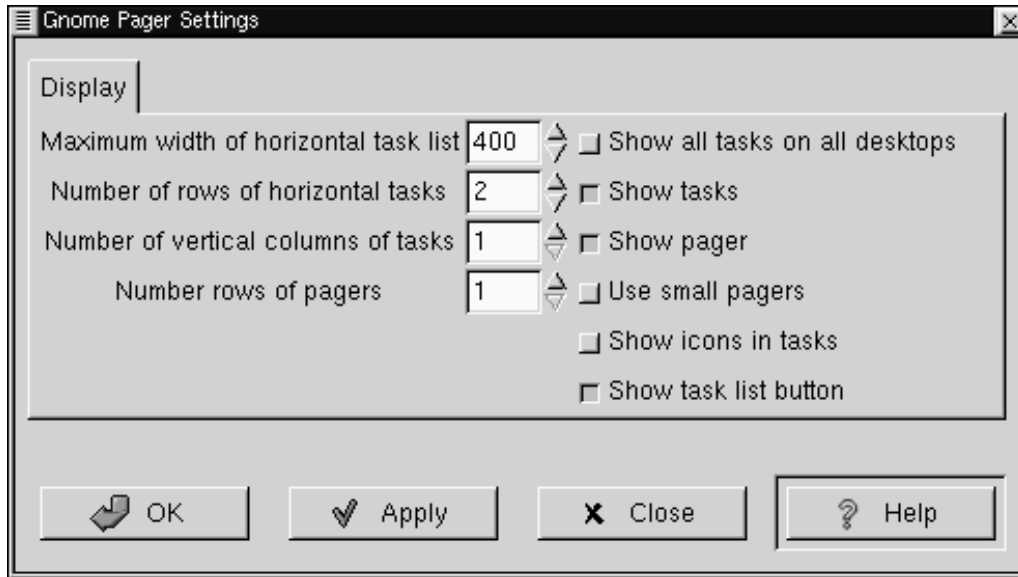
- **Pager** -

   - **Show Pager** - This allows you to show the pager or not. Not showing the Pager will allow you to just use the tasklist.

   - **Place pagers after tasklist** - This allows you to switch the location of the pager and the tasklist from left to right and top to bottom depending on the orientation of your Panel.

   - **Width of small pagers** - If you have the small pagers on you can define the width here.

   - **Height of small pagers** - If you have the small pagers on you can define the height here.

   - **Width of large pagers** - If you have the large pagers on you can define the width here.

- **Height of large pagers** - If you have the large pagers on you can define the height here.

- **Tasklist** -

  - **Show task list button** - This toggles the tasklist button. The task list button is the small button with the up arrow which brings up a list of all tasks on all desktops when pressed.

  - **Show task list** - You can choose to show the task list or use the pager alone.

  - **Show button icons** - This will toggle icons on the task list buttons. Turning them off may save you some space.

  - **Which tasks to show** - this area lets you define which applications will show up in the task list and when. You can choose to have all tasks, normal tasks only, minimized tasks only, all tasks on all desktops, or all minimized tasks on all desktops.

  - **Geometry** - this area allows you to change the geometry of the tasklist. You can have the tasklist always display at minimum size or change the width of the horizontal and vertical task lists. You can also specify the number of rows of horizontal tasks and the number of vertical columns of tasks.

19.6.5 Quicklaunch Applet

The Quicklaunch Applet is a small applet that gives you a repository to place launchers. The Quicklaunch applet holds the application launchers you wish to have and allows you to click on them to launch the applications. Drag and drop functionality makes the set up of the Quicklaunch applet very easy and quick.

Figure 19–12   The Quicklaunch Applet
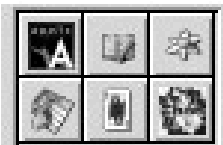
To add a launcher to the Quicklaunch applet you must already have the launcher set up either in the Main Menu, on your Panel, or on your desktop. Once you have

the launcher you can drag it onto the Quicklaunch applet to create a small launcher button. To launch the application simply press the launcher button.

To drag a launcher from your Main Menu click once on the Main Menu to open it, and click once on the sub menu to open it. Now move the mouse to the application desired, press (but do not release) the left mouse button, drag the application to the Quicklaunch applet and release the button. To drag a launcher from either the Panel or your desktop, move the mouse to the launcher, press and hold the left mouse button, drag the application to the Quicklaunch applet, and release the button.

Once the launcher is in the Quicklaunch applet you can right mouse click on it and select **Properties** to change any properties associated with the launcher. The dialog that is launched is the standard GNOME launcher properties dialog, which you can read more about in Section 12.4.1, *Adding Application Launchers*.

20   GNOME CD Player

20.1 Introduction

The GNOME CD Player (gtcd) is a GNOME enabled application that is preloaded with GNOME. This is a simple CD Player which allows you to listen to Compact Discs on your PC.

20.2 Using the GNOME CD Player

The GNOME CD Player will be available to you in the Main Menu in the Audio menu and can also be invoked in the command line with $gtcd.

Figure 20–1   The GNOME CD Player

IMPORTANT

You must have the correct access rights to your CD-ROM drive for this application to be successful. Some systems, will normally grant you the necessary rights automatically when you log into the console. The mechanism that does this is called the pam_console. If your system doesn't give you the necessary rights to the CD-ROM, then you will need to be given those rights. If you have the root password type the following in a terminal window.

```
$ su
$ Password: [type  in root password]
$ chmod a+r /dev/cdrom
$ exit
```

If your CD-ROM is located somewhere other than /dev/cdrom make sure you change it in the commands above.

The GNOME CD Player works like any CD Player with common buttons such as Play, Stop, Pause, etc. Plus a track selector button that displays the track titles in a drop down menu. You have access to change various properties by pressing the **Preferences** button. This will bring up the GNOME CD Player Preferences dialog.

There are three tabs in the GNOME CD Player Preferences dialog: Preferences, Key-bindings, and CDDB Settings.

Figure 20–2 The GNOME CD Player Properties
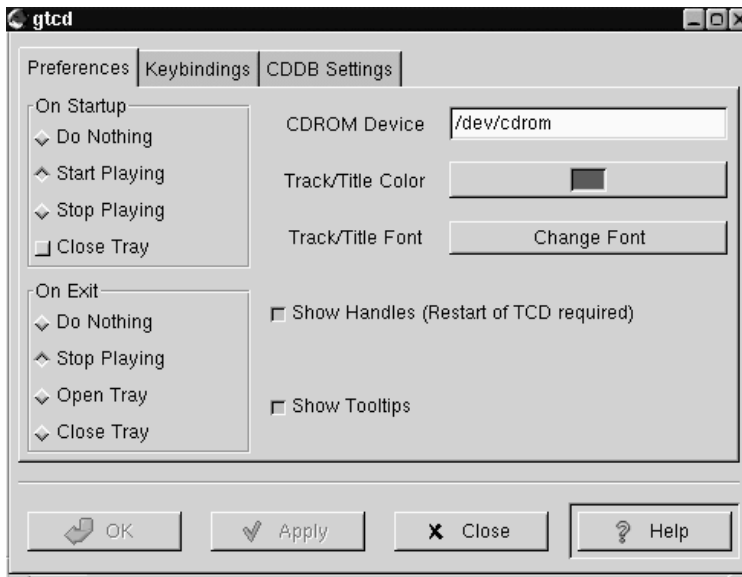
- Preferences Tab - In this dialog you can specify:

  What you would like the GNOME CD Player to do when first started, and when exited.

  The location of your CD-ROM on your system. This is usually /dev/cdrom.

  The Color to display the Track and CD Title.

  The Font to display the Track and CD Title.

  Whether you would like handles on the title window which will allow you to drag the title window off of the CD Player to float on the desktop.

  Whether tooltips are enabled when your mouse is over the buttons.

- Keybindings Tab - In this tab you can change the key bindings associated with the GNOME CD Player. These keybindings allow you to use the Player without using your mouse. If you want to change one of the bindings, select it with your

mouse and type the new key in the `Click here to change` text box. Press Apply to save the changes

• CDDB Settings Tab - CDDB stands for CD Database and is a huge global database of CD information. Each CD has an identity, which the CD Player can read. If you are connected to the internet, it will then search a CDDB server for that CD identity and return any information it has on it. This usually includes CD Title, Artist, and track titles. It can also include notes and lyric information. Once this data is retrieved, the GNOME CD Player will store the information on your hard drive for future access. In the CDDB Setting tab you can change the CDDB server and edit your local CDDB database. You can read more about CDDB by visiting the CDDB Website.

Another feature in the GNOME CD Player is the Track Editor. The track editor can be launched by the **Track Editor** button on the main GNOME CD Player window. The Track Editor allows you to edit the CD track information in case it is incorrect or there was no CDDB entry for your CD. You can also check the status of the CDDB information by pressing the **CDDB Status** button at the bottom of the Track Editor. This will show you what messages, if any, were returned from the CDDB server.

Figure 20–3   The GNOME CD Player Track Editor

```
┌─────────────────────────────────┐
│ ⌐ Track Editor          _ □ ×   │
│ Disc Information (45:54 minutes) │
│           Artist / Title         │
│ ┌─────────────────────────────┐ │
│ │ Television / Marquee Moon   │ │
│ └─────────────────────────────┘ │
│ Track Information                │
│ ┌─────────────────────────────┐ │
│ │                             │ │
│ └─────────────────────────────┘ │
│ ┌───┬──────┬──────────────────┐ │
│ │Trk│ Time │ Title            │ │
│ ├───┴──────┴──────────────────┤ │
│ │ 1   3:54  See No Evil        │ │
│ │ 2   3:51  Venus              │ │
│ │ 3   4:45  Friction           │ │
│ │ 4  10:40  Marquee Moon       │ │
│ │ 5   5:06  Elevation          │ │
│ │ 6   5:35  Guiding Light      │ │
│ │ 7   5:02  Prove It           │ │
│ │ 8   6:57  Torn Curtain       │ │
│ │                             │ │
│ ├─────────────────────────────┤ │
│ │        CDDB Status          │ │
│ ├──────────────┬──────────────┤ │
│ │   🐾 OK      │   ✖ Cancel   │ │
│ └──────────────┴──────────────┘ │
└─────────────────────────────────┘
```

21   The GNOME Calendar

21.1 Introduction

The GNOME Calendar is a simple calendar applications which can be quite useful in your daily work. The Calendar, like much of GNOME, is in its infancy and while it is a very useful application, there is some very impressive functionality still to come like network shared calendar usage, and syncing capability with Palm Pilots and other hand held PDAs. The Calendar can be launched from the **Applications** submenu of the Main Menu.

The Calendar is broken up into four main tabs: Day, Week, Month, and Year. Each view allows you to view the respective time period but will also allow you to schedule appointments from any one of the tabs.

Figure 21–1   The GNOME Calendar

21.2 Setting Up the GNOME Calendar

The GNOME Calendar can be used in the state which it is shipped but there are tools available to make sure it is setup the way you would like it to be. The first place you should go is the GNOME Calendar Preferences. You can find the Preferences dialog by selecting the **Preferences** menu item from the **Settings** menu.

Figure 21–2   The GNOME Calendar Preferences
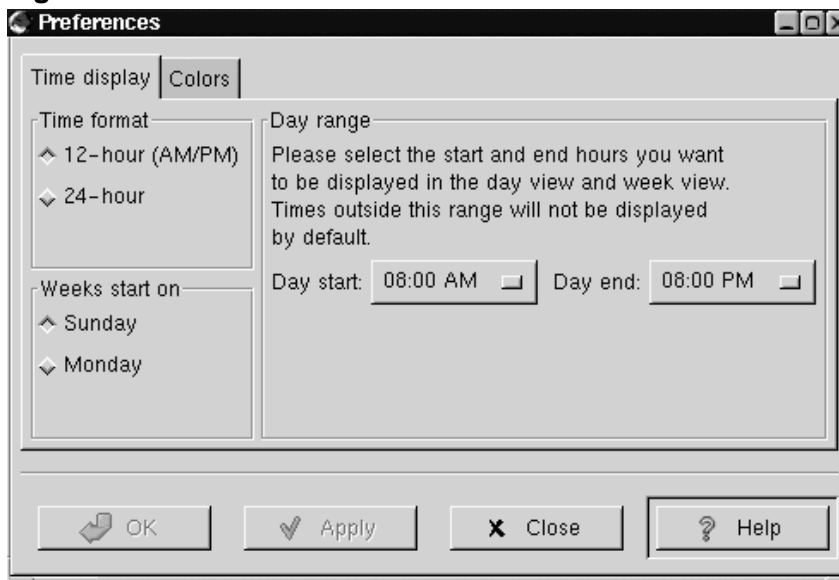
The Calendar Preferences dialog is broken up into three tabs, Time display, Colors, and Todo List

Time - This tab is broken up into three main sections: Time format, Weeks start on, and Day range.

Time format - This allows you to choose between a 12 or 24 hour format clock.

Weeks start on - This allows you to define what day your week will start on, Sunday or Monday. This will affect how the calendar is laid out in the Day and Week views.

Day range - This section lets you choose what time your days will start and end. Any hours outside the range selected will be shaded on your Day view.

Colors - This tab allows you to change the default colors used in the Calendar. There are seven color choices you can customize: Outline, Headings, Empty days, Appointments, Highlighted day, Day numbers, and Current day's number. Each choice has a small color selector box next to it. When you press this box you will be given a color selector dialog in which you can choose the color you want. Once you have selected a color the small sample calendar on the right side of the tab will preview your choice.

Todo List - This tab allows you to define which columns will be shown in the Todo List in the Day View. You can set the Summary, Due Date, and Priority columns.

Once you have made the changes to the Calendar Preferences you can press the **Apply** button to apply them.

21.3 Using the GNOME Calendar

Using the GNOME Calendar is quite simple and most tasks can be performed from any of the major views, Day, Week, Month or Year. Probably one of the most important features to remember is that at any time you can right mouse click on a particular day and add a new appointment. There are many other features which will be described below in each of the major views.

21.3.1 The Day View

The Day tab is probably the most useful view in the GNOME Calendar as it acts just as a day timer would. On the left of the tab is the hour listing for the current day. The light grey coloring in the hour list separates the work hours from the non-work hours. If you would like to change the work hours displayed in light grey you can do so in Section 21.2, *Setting Up the GNOME Calendar*.

Figure 21–3   The Day View

TIP

One tip for adding a new appointment in the Day View is to select a few hours in the hours list by clicking and dragging your mouse down the hours list. Once the correct range of time has been selected you can press Enter and type in the appointment. This will allow you to skip the Create New Appointment dialog.

Next to the hours listing in the top right hand corner is a small full month calendar. You can change the month or year of the small month calendar by pressing the forward and backward arrows on the top. You can use the small month calendar to navigate the days as well. Double clicking on any day in the small month calendar will move the current day view to that particular day.

Below the small month calendar is your Todo list. The Todo list is a simple list where you can keep all your tasks on hand. To add an item to the Todo list press the **Add** button. This will launch a small editing box where you can type in the item. Once you have entered an item in the To-do list you can use the **Edit** and **Delete** to manage you items. The To-do items are available no matter which days are displayed in the Day View and can only be deleted with the **Delete** button.

21.3.2 The Week View

The Week View shows the current week with detailed descriptions of your appoint-ments. If you would like to add an appointment for any of the days in the week view, you can right mouse click on the day and select the **New appointment** menu item from the **pop-up** menu. You can also use the week view to navigate to particular days in the Day View. Double-click on any day in the Week View and you will go to that day in the Day View.

Figure 21–4   The Week View

In the lower left corner of the Week View there is a small month calendar. You can change the month or year of the small month calendar by pressing the forward and backward arrows on the top. You can use the small month calendar to navigate the days as well. Double clicking on any day in the small month calendar will move the current week view to that particular week.

21.3.3 The Month View

The Month View shows the entire month with brief detailed descriptions of your appointments. The Month View makes use of the customized colors available in the Calendar. You can read about how to set these colors in Section 21.2, *Setting Up the GNOME Calendar*. For any day with a brief description of an appointment you can click on the day to display a detailed description of the appointment in a pop-up window. If you would like to add an appointment to a day in the Month View you can right mouse click on any day and select the **New Appointment in this day** item from the **pop-up** menu. You can use the items in the **pop-up** menu to navigate in the Day, Week, and Year views by selecting either **Jump to this day**, **Jump to this week**, or **Jump to this Year**.

Figure 21–5   The Month View

21.3.4 The Year View

The Year View shows you the entire year with no descriptions of appointments. Like the Month View, the Year View makes use of the customized colors available in the Calendar. You can read about how to set these colors in Section 21.2, *Setting Up the GNOME Calendar*. If you have an appointment on a day you can click on that day and a description of the appointment will be displayed in the pop-up window. If you would like to add an appointment to a day in the Year View you can right mouse click on any day and select the **New Appointment in this day** item from the **pop-up** menu. You can use the items in the **pop-up** menu to navigate in the Day, Week, and Month views by selecting either **Jump to this day**, **Jump to this week**, or **Jump to this Month**.

Figure 21–6   The Year View

21.4 Making a New Appointment

There are many methods for making a new appointment in the GNOME Calendar, the easiest is by pressing the **New** button on the button bar. Whenever you make a new appointment you will launch the Create New Appointment dialog that allows you to set the properties of that appointment. The Create New Appointment dialog is broken into two different tabs, the General and the Recurrence tabs.

General - The General tab is the area in which you specify the time of the appointment and set reminders for yourself. There are four different areas on the General tab: Summary, Time, Alarms, and Classification.

•   Summary - The Summary box allows you to type a description of the appoint-ment. Keep in mind that only a portion of this description will be available in the Week and Month Views.

- Time - The Time area allows you to set the time range for the appointment by selecting the date and hours. To the right of the start and end days there is a small selection box named Calendar. This will bring up a small Calendar when pressed. You can select the start and end date in the small calendar. To the right of the start and end hours there is a small button that will display the hours of the day when pressed. Each hour in the list will have a sub menu displaying each quarter hour so you can select them. The hours are restricted to those included in the Day Range.

- Alarms - The Alarms area allows you to set up an alarm to remind you of an appointment. There are four different types of alarms you can use to remind yourself of appointments: Display, Audio, Program, and Mail. The Display alarm will display a message on your screen at the time you set. The Audio alarm will play an audio file at the time you set. The Program alarm will run a program you specify at the time you set. The Mail alarm will send email to the user specified at the time you set.

Recurrence - The Recurrence tab allows you to specify how often an appointment should recur, if at all. The first property you should set if you want a recurring appointment is the Recurrence rule. You can choose among None, Daily, Weekly, Monthly, and Yearly. For each selection you can adjust the recurrence properties for your selection. In the Ending date area you can set a rule which will stop the recurrence of your appointment or allow it to repeat forever. In the Exceptions area you can make exceptions to the recurring appointment by double clicking the date and pressing the **Add** button.

GNU General Public License

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

675 Mass Ave, Cambridge, MA 02139, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that re-distributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

1.   This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

2.   You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

3.   You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a.   You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b.   You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c.   If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

4.   You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

   a.   Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

   b.   Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

   c.   Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

5. You may not copy, modify, sub-license, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sub-license or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

6. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

7. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

8. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

    If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

9.   If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

10. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

11. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

12.   BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

13.   IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Part III   Appendixes

A   Glossary

Alpha

A RISC (*R*educed *I*nstruction *S*et *C*omputer) architecture developed by Digital Equipment Corporation.

ATAPI

An abbreviation for *AT A*ttachement *P*acket *I*nterface. ATAPI is the protocol by which CD-ROM drives communicate with a computer system over an IDE interface.

Binary

Although the base two-numbering system used by computers is known as binary, the word often refers to the executable form of a program. Contrast with "source code."

BIOS

An abbreviation for *B*asic *I*nput/*O*utput *S*ystem. On PC-compatible systems, the BIOS is used to perform all necessary functions to properly initialize the system's hardware when power is first applied. The BIOS also controls the boot process, provides low-level input/output routines (hence its name) and (usually) allows the user to modify details of the system's hardware configuration.

Boot

Short for "bootstrap." The process by which a computer starts running an operating system when power is applied.

Boot Disk

A diskette used to start many Red Hat Linux installations.

Bootstrap

See "Boot."

CISC

An abbreviation for *C*omplex *I*nstruction *S*et *C*omputer. A design philosophy for computers whereby the processor is designed to execute a relatively large number of different instructions, each taking a different amount of time to execute (depending on the complexity of the instruction). Contrast with RISC.

CMOS

Originally an abbreviation for *C*omplementary *M*etal *O*xide *S*emiconductor -- a semiconductor technology used in many integrated circuits. Now often used to describe the low-level hardware that contains a personal computer's BIOS setting, and the computer's hardware clock.

Cylinder

When referring to disk drives, the number of different positions the disk drive's read/write heads can take over the unit's disk platters. When viewed from above the platters, each head position describes an imaginary circle of different diameters on the platter's surface, but when viewed from the side, these circles can be thought of as a series of cylinders nested within each other, hence the term. See also Geometry.

Daemon

A daemon is a program that runs, without human intervention, to accomplish a given task. For example, `lpd` is a daemon that controls the flow of print jobs to a printer.

Dependencies

When referring to packages, dependencies are requirements that exist between packages. For example, package `foo` may require files that are installed by package `bar`. In this example, `bar` must be installed, or else `foo` will have unresolved dependencies. RPM will not normally allow packages with unresolved dependencies to be installed.

Device Driver

Software that controls a device that is connected to, or part of, a computer.

Disk Drive

See Hard Disk.

Disk Druid

Disk Druid is a component of the Red Hat Linux installation program that is used to partition disk drives during the installation process.

Diskette

A small mass storage device in a removable cartridge, meant to be read/written to, in a compatible drive.

Distribution

An operating system (usually Linux) that has been packaged so as to be easily installed.

Domain Name

A domain name is used to identify computers as belonging to a particular organization. Domain names are hierarchical in nature, with each level in the hierarchy being separated from other levels with a period (pronounced "dot"). For example, Foo Incorporated's Finance department might use the domain name "finance.foo.com."

Driver

See Device Driver.

Dual Boot

The act of configuring a computer system to boot more than one operating system. The name is something of a misnomer, as it is possible to boot more than the two operating systems the word "dual" implies.

EIDE

An abbreviation for *E*nhanced *I*ntegrated *D*rive *E*lectronics, which is a newer version of the IDE interface standard and another term for a particular implementation for IDE interfaces. EIDE makes larger and faster disk drives possible; most systems sold today use EIDE.

Errata

Errata is Latin for "Ooops."[1] When software is found to have bugs, quite often the software is fixed, and released as errata. Red Hat Linux is no exception to the rule; we have an Errata Web page at http://www.redhat.com/errata.

Extended Partition

A segment of a disk drive that contains other partitions. See Partition.

FAQ

An abbreviation for Frequently Asked Questions. Linux information is often presented in the form of lists of questions and answers called FAQs.

fdisk

fdisk is a utility program that is used to create, delete or modify partitions on a disk drive.

Filesystem

A filesystem is the method by which information is stored on disk drives. Different operating systems normally use different filesystems, making it difficult to share the contents of a disk drive between two operating systems. However, Linux supports multiple filesystems, making it possible, for example, to read/write a partition dedicated to Windows.

Floppy

A somewhat historical term for a small mass storage device in a removable cartridge, meant to be read/written to in a compatible drive. See "diskette."

Formatting

The act of writing a filesystem on a disk drive.

---

[1] Well, it should be…

FQDN

An abbreviation for *F*ully *Q*ualified *D*omain *N*ame. An FQDN is the human-readable name that includes a computer's hostname and associated domain name. For example, given a hostname of "foo," and a domain name of "bar.com," the FQDN would be "foo.bar.com."

FTP

An abbreviation for *F*ile *T*ransfer *P*rotocol. Also the name of a program that, as the name implies, permits the copying of files from one system on a network to another.

Gateway

In networking terms, refers to a device that connects one or more computers on a network to other networks. The device may be specialized hardware (such as a router), or may be a general-purpose computer system configured to act as a gateway.

Geometry

When referring to disk drives, the physical characteristics of the disk drive's internal organization. Note that a disk drive may report a "logical geometry" that is different from its "physical geometry," normally to get around BIOS-related limitations. See also Cylinder, Head and Sector.

GID

Short for *G*roup *ID*. The means by which a user's membership in a group is identified to various parts of Red Hat Linux. GIDs are numeric, although human-readable names are stored in the /etc/group file.

Group

Groups are a way of assigning specific access rights to certain classes of users. For example, all users working on Project X could be added to group xproj. System resources (such as disk space) devoted to Project X could then be configured to permit only members of xproj full access.

Hard Disk

A hard disk contains rotating magnetic media (in the shape of disks) that spin rapidly. Small heads float over the surface of each disk, and are used to write to and read from the disk as it rotates.

Head

When referring to disk drives, the number of read/write heads within a disk drive. For each platter in a disk drive, there are normally two heads for each platter -- one for each surface -- although one surface may go unused. See also Geometry.

Hostname

A hostname is a human-readable string of characters used to identify a particular computer system.

I18n

See Internationalization.

IDE

An abbreviation for *I*ntegrated *D*rive *E*lectronics, which is the name of a standard interface used to connect primarily disk and CD-ROM drives to a computer system. See also "EIDE" and "ATAPI."

Intel

Company responsible for producing the microprocessors that most commonly appear in PC-compatible personal computers. These processors include the 80386, 80486, Pentium, Pentium Pro, and Pentium II.

Internationalization

The practice of designing and writing programs that can be easily configured to interact with the user in more than one language. Often referred to as "i18n," due to the number of letters between the starting "i" and the ending "n."

IP Address

IP addresses are the method by which individual computer systems (or from a more strictly accurate interpretation, the network interfaces on those computer systems) are identified on a TCP/IP network. All IP addresses consist of four number blocks, each ranging from 0 to 255, and separated by periods.

Kernel

The central part of an operating system upon which the rest of the operating system is based.

Library

When speaking of computers, refers to a collection of routines that perform operations which are commonly required by programs. Libraries may be shared, meaning that the library routines reside in a file separate from the programs that use them. Library routines may also be "statically linked" to a program, meaning that copies of the library routines required by that program are physically added to the program. Such statically linked binaries do not require the existence of any library files in order to execute. Programs linked against shared libraries will not execute unless the required libraries have been installed.

LILO

A commonly-used bootstrap loader for Linux systems based on an Intel-compatible processor.

Linus Torvalds

Created Linux in 1991 while a university student.

Linuxconf

A versatile system configuration program written by Jacques Gelinas. Linuxconf provides a menu-based approach to system configuration via several different user interfaces.

Linux

A full-featured, robust, freely-available operating system originally developed by Linus Torvalds.

Logical Partition

A partition that exists within an extended partition. See also "partition" and "extended partition."

Master Boot Record

The master boot record (or MBR) is a section of a disk drive's storage space that is set aside for the purpose of saving information necessary to begin the bootstrap process on a personal computer.

MBR

See "Master Boot Record."

Memory

When referring to computers, memory (in general) is any hardware that can store data for later retrieval. In this context, memory usually specifically refers to RAM.

MILO

A commonly-used bootstrap loader for Linux systems based on the Alpha processor.

Module

In Linux, a module is a collection of routines that perform a system-level function, and may be dynamically loaded and unloaded from the running kernel as required. Often containing device drivers, modules are tightly bound to the version of the kernel; most modules built from one version of a kernel will not load properly on a system running another kernel version.

Mount

The act of making a filesystem accessible to a system's users.

Mount Point

The directory under which a filesystem is accessible after being mounted.

Nameserver

In TCP/IP networking terms, a nameserver is a computer that can translate a human-readable name (such as "foo.bar.com") into a numeric address (such as "10.0.2.14").

Netmask

A netmask is a set of four number blocks separated by periods. Each number is normally represented as the decimal equivalent of an eight-bit binary number, which means that each number may take any value between 0 (all eight bits cleared) and 255 (all eight bits set). Every IP address consist of two parts (the network address and the host number). The netmask is used to determine the size of these two parts. The positions of the bits that are set in the netmask are considered to represent the space reserved for the network address, while the bits that are cleared are considered to represent the space set aside for the host number.

NFS

An abbreviation for *N*etwork *F*ile *S*ystem, NFS is a method of making the filesystem on a remote system accessible on the local system. From a user's perspective, an NFS-mounted filesystem is indistinguishable from a filesystem on a directly-attached disk drive.

Operating System

A collection of software that controls various resources of a computer.

Packages

Files that contain software, and written in a particular format that enables the software to be easily installed and removed.

PAM

An acronym for *P*luggable *A*uthentication *M*odules. PAM is an authentication system that controls access to Red Hat Linux.

Partition

A segment of a disk drive's storage space that can be accessed as if it was a complete disk drive.

Partition Table

The partition table is a section of a disk drive's storage space set aside to define the partitions that exist on that disk drive.

Partition Type

Partitions contain a field that is used to define the type of filesystem the partition is expected to contain. The partition type is actually a number, although many times the partition type is referred to by name. For example, the "Linux Native" partition type is 82. Note that this number is hexadecimal.

PC Card

See PCMCIA.

PCMCIA

Acronym for *P*ersonal *C*omputer *M*emory *C*ard *I*nternational *A*ssociation. This organization produced a series of standards that define the physical, electrical and software characteristics of small, credit card-sized devices that can contain memory, modems, network adapters and more. Also known as PC Cards, these devices are mainly used in laptop computers (although some desktop systems can use PCMCIA cards, too).

PCMCIA Support Diskette

A diskette required for Red Hat Linux installations that require the use of a PCMCIA device during the install.

Permissions

The set of identifiers that control access to files. Permissions consist of three fields: user, group, and world. The user field controls access by the user owning the file, while the group field controls access by anyone matching the file's group specification. As the name implies, the world field controls access by

everyone else. Each field contains the same set of bits that specify operations that may or may not be performed, such as reading, writing and executing.

PLIP

An abbreviation for *P*arallel *L*ine *I*nternet *P*rotocol. PLIP is a protocol that permits TCP/IP communication over a computer's parallel port using a specially-designed cable.

POSIX

A somewhat mangled abbreviation for *P*ortable *O*perating *S*ystem *I*nterface. A set of standards that grew out of the UNIX operating system.

Process

A process (in somewhat simplistic terms) is one instance of a running program on a Linux system.

PS/2 Mouse

A PS/2 mouse gets its name from the original computer in which this type of mouse was first used -- the IBM PS/2. A PS/2 mouse can be easily identified by the small, round connector at the end of its cable.

RAM

An acronym for *R*andom *A*ccess *M*emory. RAM is used to hold programs while they are being executed, and data while it is being processed. RAM is also volatile, meaning that information written to RAM will disappear when the computer's power is turned off.

Reboot

To restart the boot process. See also "Boot."

Red Hat, Incorporated

A North Carolina software company. Produces and markets sofware for the Linux operating system, including Red Hat Linux.

Rescue Diskette

A diskette containing a rudimentary system environment. As the name implies, a rescue diskette is normally used in an attempt to "rescue" an ailing system from the necessity of re-installing the entire operating system.

RISC

An abbreviation for *R*educed *I*nstruction *S*et *C*omputer. A design philosophy for computers whereby the processor is optimized to execute a relatively small number of different instructions in a predictably small amount of time. Contrast with CISC.

ROM

An abbreviation for *R*ead *O*nly *M*emory. ROM is used to hold programs and data that must survive when the computer is turned off. Because ROM is non-volatile; data in ROM will remain unchanged the next time the computer is turned back on. As the name implies, data cannot be easily written to ROM; depending on the technology used in the ROM, writing may require special hardware, or may be impossible. A computer's BIOS may be stored in ROM.

Root

The name of the login account given full and complete access to all system resources. Also used to describe the directory named "/"as in, "the root directory."

RPM

An abbreviation that stands for *R*ed Hat *P*ackage *M*anager. RPM is also the name of a program that enables the installation, upgrading and removal of packages.

SCSI

An abbreviation for *S*mall *C*omputer *S*ystem *I*nterface, SCSI is a standard interface for connecting a wide variety of devices to a computer. Although the most popular SCSI devices are disk drives, SCSI tape drives and scanners are also common.

Sector

When referring to disk drives, the number of fixed-size (normally 512 byte) areas that can be accessed by one of the disk drive's read/write heads, in one rotation of the disk, without that head changing position. See Also Geometry.

Serial Mouse

A serial mouse is a mouse that is designed to be connected to a computer's serial port. A serial mouse can be easily identified by the rectangular-shaped connector at the end of its cable.

setgid

A system call that can be used to set the GID of a process. Programs can be written using setgid such that they can assume the group ID of any group on the system.

setuid

A system call that can be used to set the UID of a process. Programs can be written using setuid such that they can assume the user ID of any process on the system. This is considered a possible security problem if a program is "setuid root."

Shadow Password

Normally, each user's password is stored, encrypted, in the file /etc/passwd. This file must be readable by all users so that certain system functions will operate correctly. However, this means that copies of user's encrypted passwords are easily obtained, making it possible to run an automated password-guessing program against them. Shadow passwords, on the other hand, store the encrypted passwords in a separate highly-protected file, making it much more difficult to crack passwords.

SILO

A commonly-used bootstrap loader for Linux systems based on the SPARC processor.

SLIP

An acronym for *S*erial *L*ine *I*nternet *P*rotocol. SLIP is a protocol that permits TCP/IP communication over serial line (typically over a dial-up modem connection).

SMB

Short for *S*erver *M*essage *B*lock, SMB is the communications protocol used by Windows-based operating systems to support sharing of resources across a network.

source code

The human-readable form of instructions that comprise a program. Also known as "sources." Without a program's source code, it is very difficult to modify the program.

SPARC

A RISC (*R*educed *I*nstruction *S*et *C*omputer) architecture developed by Sun Microsystems.

Swap

Also known as "swap space." When a program requires more memory than is physically available in the computer, currently-unused information can be written to a temporary buffer on the hard disk, called swap, thereby freeing memory. Some operating systems support swapping to a specific file, but Linux normally swaps to a dedicated swap partition. A misnomer, the term swap in Linux is used to define damand paging.

System Call

A system call is a routine that accomplishes a system-level function on behalf of a process.

TCP/IP

An abbreviation for *T*ransmission *C*ontrol *P*rotocol/*I*nternet *P*rotocol, TCP/IP is the name given to the networking standard commonly used on the Internet today.

Torvalds, Linus

See Linus Torvalds.

UID

Short for *U*ser *ID*. The means by which a user is identified to various parts of Red Hat Linux. UIDs are numeric, although human-readable names are stored in the /etc/passwd file.

UNIX

A set of Linux-like operating systems that grew out of an original version written by some guys at a phone company.[2]

Unmount

The act of revoking access to a filesystem. (Note that the program that unmounts filesystems is called umount.)

Virtual Console

Virtual consoles provides multiple "screens" on which a user may log in and run programs. One screen is displayed on the computer's monitor at any given time; a key sequence is used to switch between virtual consoles.

Widget

A standardized on-screen representation of a control that may be manipulated by the user. Scroll bars, buttons, and text boxes are all examples of widgets.

X Window System

Also known as "X," this graphical user interface provides the well-known "windows on a desktop" metaphor common to most computer systems today. Under X, application programs act as clients, accessing the X server, which manages all screen activity. In addition, client applications may be on a different system than the X server, permitting the remote display of the applications graphical user interface.

---

[2] Just kidding -- thank you Ken Thompson and Dennis Ritchie of Bell Telephone Laboratories for your inspired operating system design!

XFree86

A free implementation of the X Window System.

Index