

Red Hat Linux 7.0

The Official Red Hat Linux Getting Started Guide

ISBN: 1-58569-076-7

Red Hat, Inc.
2600 Meridian Parkway Durham NC 27713 US 919-547-0012 1-888-733-4281 919-547-0024
docs@redhat.com 13588 Research Triangle Park NC 27709

© 2000 Red Hat, Inc.

GSG(EN)-7.0-Print-RHI (2000-08-09T14:34-0400)

Red Hat is a registered trademark and the Red Hat Shadow Man logo, RPM, the RPM logo, and Glint are trademarks of Red Hat, Inc.

Linux is a registered trademark of Linus Torvalds.

Motif and UNIX are registered trademarks of The Open Group.

Alpha is a trademark of Digital Equipment Corporation.

SPARC is a registered trademark of SPARC International, Inc. Products bearing the SPARC trademark are based on an architecture developed by Sun Microsystems, Inc.

Netscape is a registered trademark of Netscape Communications Corporation in the United States and other countries.

TrueType is a registered trademark of Apple Computer, Inc.

Windows is a registered trademark of Microsoft Corporation.

All other trademarks and copyrights referred to are the property of their respective owners.

Copyright © 2000 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.

Distribution of the work or derivative of the work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

Printed in the United States, Ireland, and Japan

This manual is dedicated to Paul Gallagher, creator of the *Official Red Hat Linux Getting Started Guide* and first Club Docs editor. Paul, thanks for all you've done to make Red Hat Linux documentation what it is today. We wish you every *Success!*

Contents

Red Hat Linux 7.0

Introduction	ix
Who Should Read This Book?.....	ix
A Peek Inside	xi
A Note About Environments	xi
Work With Your User Account	xiii
Where to Find More Info.....	xiv
More to Come	xv
Sign Up for Support.....	xv
Part I The Basics	17
Chapter 1 Getting Started	19
1.1 Log In, Please.....	19
Chapter 2 GNOME in 60 Seconds (or so)	31
2.1 An Overview	31
2.2 The Desktop	33
2.3 Folders.....	35
2.4 The Panel	37
2.5 Main Menu Button	39
2.6 Finding Help.....	40
2.7 Multiple Desktops.....	41
2.8 Switching Environments	43
Chapter 3 KDE in 60 Seconds (or so)	45
3.1 An Overview	45
3.2 The Desktop	49
3.3 Folders.....	52
3.4 The Panel	54

3.5	The Main Menu K Icon.....	58
3.6	The Taskbar	59
3.7	Finding Help.....	60
3.8	Switching Environments	63
Chapter 4	Logging Out and Shutting Down	65
4.1	Quitting from GNOME.....	65
4.2	Quitting from KDE	68
4.3	Shutting Down at the Shell Prompt	69
Chapter 5	Setting Up Your Printer	71
5.1	Using Printtool	72
5.2	Printing Utilities in KDE and GNOME	82
Chapter 6	Getting Online	91
6.1	Using RP3 to Connect	92
6.2	Connecting with Kppp.....	108
Chapter 7	The Web, E-Mail and News in GNOME and KDE	129
7.1	Using Navigator in GNOME and KDE.....	130
7.2	E-Mail and News Settings for Messenger.....	138
7.3	Reading and Composing Mail in Messenger	147
7.4	Using Messenger to Read News.....	155
Chapter 8	Customizing Your Desktop in GNOME and KDE	161
8.1	Customizing with the GNOME Control Center	161
8.2	Customizing with the KDE Control Center	165
Chapter 9	A Short Introduction to the GIMP	169
9.1	GIMP Basics	170
9.2	Using the GIMP to Create a Web Page	187
9.3	Where to Find More Information	212

Part II Managing Files and Directories	215
Chapter 10 File Managers in GNOME and KDE	217
10.1 The GNOME File Manager	217
10.2 The KDE File Manager.....	228
Chapter 11 An Introduction to Linuxconf	237
11.1 Creating a New Account in Linuxconf	237
Chapter 12 Using Gnome-RPM and RPM	245
12.1 The benefits of RPM	245
12.2 Using Gnome-RPM.....	246
12.3 Installing New Packages	249
12.4 Basic Configuration.....	251
12.5 Package Manipulation.....	256
12.6 RPM from the Shell Prompt	260
12.7 Checking a Package's "Signature".....	267
Chapter 13 Getting Started with GnuPG	273
13.1 Introducing GnuPG	273
Part III The Shell Prompt	285
Chapter 14 Shell Prompt Basics	287
14.1 Why Use a Shell Prompt?	287
14.2 Pulling Yourself Up by the Boot	287
14.3 A Good "Man" Is Easy to Find	289
Chapter 15 You Are Here	293
15.1 Finding Yourself with pwd.....	293
15.2 Getting from Here to There: cd	294
15.3 Looking Around with ls.....	302

15.4	A Larger Picture of the Filesystem.....	311
15.5	"Washing" the Window.....	315
15.6	Using cat.....	316
15.7	Using Redirection.....	318
15.8	Appending Standard Output.....	322
15.9	Redirecting Standard Input.....	324
15.10	Pipes.....	325
15.11	Stringing Commands Together.....	327
15.12	Ownership and Permissions.....	327
15.12	Identities.....	333
15.12	Permissions.....	333
15.12	Actions.....	333
15.13	Fun with Numbers and chmod.....	336

Chapter 16 Working with Files and Directories..... 341

16.1	Shell Collecting.....	341
16.2	Locating Files and Directories.....	343
16.3	Command History and Tab Completion.....	345
16.4	Identifying and Working with File Types.....	348
16.5	Copying, Moving and Renaming Files and Directories.....	355

Part IV Q & A..... 361

Chapter 17 Q & A: Quick Answers to Common Questions .. 363

17.1	Your First Login.....	363
17.2	Using a Diskette.....	364
17.3	Switching Environments.....	366
17.4	Shutting Down.....	367
17.5	Error Message While Installing an RPM.....	368
17.6	Starting Applications.....	369
17.7	Accessing a Windows Partition.....	371
17.8	Find Some Commands Fast.....	373
17.9	Keep ls Output from Scrolling.....	374
17.10	Tips on Using Command History.....	375

17.11	Unable to Log In	376
17.12	Starting "Services"	377
17.13	Using All Available RAM	377
17.14	Configuring a Sound Card	380
17.15	Removing Red Hat Linux	383
17.16	Changing Login from Console to X at Startup	384
17.17	Configuring X	386
Part V Appendixes		393
Appendix A Glossary		395
Appendix B A Brief History of Linux and Red Hat		411
B.1	What is Linux?	411
Appendix C DOS to Linux Cheat Sheet		415
C.1	DOS and Linux Commands	415
C.2	Different Paths	417

Introduction

Welcome to the *Official Red Hat Linux Getting Started Guide*!

Just a few short years ago, a user who wanted to try out Linux had relatively few choices in the **distributions**, or types of packaged Linux operating systems, which were available. That's all changed. Now, there are dozens of Linux distributions from which you can choose.

But Red Hat Linux has remained a leading choice of Linux distributions for years. Here at Red Hat, we believe we offer the best Linux distribution on the market. We hope you'll agree that the time and money you spent to install and work with Red Hat Linux were well spent, indeed.

The *Official Red Hat Linux Getting Started Guide* is part of our effort to help you become acquainted with your new Red Hat Linux system.

Who Should Read This Book?

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 and tricks which can help you get acquainted with your desktop environment. Concentrating on tasks in GNOME and KDE, you'll see how to navigate through your system, how to manage files and directories and how to make some of the power of Red Hat Linux work for you.

If you're most familiar with a Windows environment, you'll find the DOS-to-Linux commands in Appendix C, *DOS to Linux Cheat Sheet* to be a useful guide, so you can start looking around right away.

Figure 1 A GNOME Desktop



You'll learn how to accomplish such basic first steps as connecting to the Internet, configuring your e-mail client, adding a printer and more.

More advanced topics will cover how to work with Update Agent, so your system automatically stays current with critical software updates.

Is Red Hat Linux sharing your computer with Microsoft Windows? Turn to Chapter 17, *Q & A: Quick Answers to Common Questions* to find out how you can allow Red Hat Linux to "see" into your other operating system, so you can perform certain tasks without having to reboot into the other operating system.

You'll even gain some experience at file management and system administration using the "shell prompt," the command-line interface.

In short, this guide picks up where the *Official Red Hat Linux Installation Guide* leaves off: by giving you the confidence to work with your new operating system.

A Peek Inside

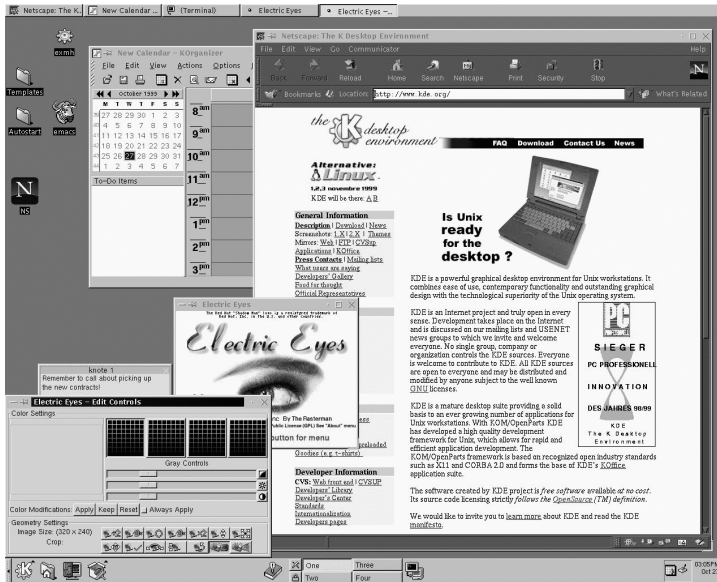
This book is divided into several parts:

- **The Basics:** When you've finished this section, you'll start to feel comfortable with your GNOME or KDE environment. You'll take your first steps to becoming productive (or just having fun and looking productive). You'll find out how to get online, set up an e-mail account and communicate with others using your Netscape Navigator Web browser.
- **Managing Files and Directories:** Beyond the basics; here, you'll learn about the GNOME and KDE file managers, working with running processes, locating files, installing and removing applications, using Update Agent and other administrative necessities.
- **The Shell Prompt:** Just beneath your graphical environment, the real power of your system awaits. You'll get a taste of this power and learn how to put it to work for you.
- **Q & A:** When you care only for answers, turn to this chapter. Many of the entries you'll find here concern answers to some of the most common new user questions, such as: how to access a floppy disk; what happens if you forget your root password (you can create a new one); how to access files on a Windows partition; and more.
- **The Glossary:** Here, you'll find quick definitions to plenty of words and phrases that can vex any new Linux user.

A Note About Environments

Your Red Hat Linux system is more than the graphical user interface, but it's a fairly safe bet that most new and intermediate users will choose to work within the graphical environment, so we'll spend most of the time demonstrating the techniques behind accomplishing tasks in two popular environments: GNOME and KDE.

Figure 2 A KDE Desktop



Much of what you can accomplish from a graphical environment can also be performed from the shell prompt. In fact, many actions you can perform from the shell prompt are faster than from the graphical environment.

In addition to the chapters covering tasks you can perform in KDE and GNOME, you'll find plenty of useful information within covering basic information from the shell prompt (as shown in Figure 3, *Using Pico at a Shell Prompt*).

Figure 3 Using Pico at a Shell Prompt

```

Terminal
File Edit Settings Help
UW PICO(tm) 3.5      File: .bashrc      Modified
# .bashrc
# User specific aliases and functions
alias rm='rm -i'
alias mv='mv -i'
alias cp='cp -i'
#####export SGML_CATALOG_FILES=/usr/lib/sgml/CATALOG
# Source global definitions
if [ -f /etc/bashrc ]; then
    . /etc/bashrc
fi
^G Get Help  ^O WriteOut  ^R Read File ^Y Prev Pg   ^K Cut Text  ^C Cur Pos
^X Exit      ^J Justify   ^W Where is  ^U Next Pg   ^U UnCut Text ^T To Spell

```

You'll also find useful tips, caution notices and related information inside.

And speaking of caution...

Work With Your User Account

Linux is a multi-tasking, multi-user operating system, so it can easily accommodate many users' needs. Even if your machine isn't attached to a network, you can share your machine with others -- each of whom will have their own settings.

Because of **user accounts**, Linux makes it easy to share your machine while preserving customized settings. Using accounts is convenient for users; it's also secure, because ordinary users don't have **permission** to damage or delete sensitive system data. Only one account, the **root account** can make changes to the system -- and to others' account information.

Using your system's root account for everyday work can be dangerous, because you can easily damage your system by accidentally deleting or modifying a sensitive file.

Most tasks in the following chapters don't require that you be logged in to your system's root account. Instead, you can accomplish almost everything from your user account -- so you won't accidentally harm your system.

Figure 4 Log In To Your User Account



Unless you're specifically told that you should be logged in as root, log in to your user account to accomplish tasks in this manual -- without risking disaster.

Where to Find More Info

Plenty of websites offer valuable information on how to use a command, but you've probably got quite a bit of documentation already on your system. In the following chapters, you'll learn:

- How to read and print **man pages** and **info pages**. Man and info pages are system documents which can help you understand what many commands are for and how to use them.
- How to use the GNOME and KDE help browsers.
- Some useful online resources you should bookmark in your browser.

- And in Appendix B, *A Brief History of Linux and Red Hat*, you'll find a brief history of Linux -- and Red Hat.

In addition, if you've got an Official Red Hat Linux boxed set, remember to take a look through the Red Hat Linux 7.0 Documentation CD for quick access to the *Official Red Hat Linux Reference Guide* and other reference material from the Linux Documentation Project.

More to Come

As Linux evolves, so will the *Official Red Hat Linux Getting Started Guide*. In future editions, expect to find more essential information to help you get the utmost from your 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)-7.0-Print-RHI (2000-08-09T14:34-0400)

That way we'll know exactly which version of the guide you have. Send your feedback to:

docs@redhat.com.

Sign Up for Support

If you have an official edition of Red Hat Linux 7.0, please remember to sign up for the benefits you're entitled to as a Red Hat customer.

You can receive 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 7.0 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** -- Use Red Hat's Update Agent to automatically retrieve and install system and security updates as well as 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.

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

The Red Hat Documentation Team

Part I The Basics

1 Getting Started

Like many new users, learning how to work on your Red Hat Linux system can be both exciting and a little unnerving. To begin your journey, you'll have to log in. When you log in, you're basically "introducing yourself" to the system.

Linux is case sensitive

Like UNIX, Linux is case sensitive. That means that typing `root` refers to a different account than `Root`, as far as Linux is concerned -- the lowercase `root` is the name of the **root login**, or system administrator.

When you installed Red Hat Linux, you had the opportunity to install the **X Window System** -- also simply called **X** -- which is the graphical environment. You were also asked whether you wanted to use a graphical screen, rather than a **console** to log in. Although our emphasis throughout this book will be on navigation and productivity using **X**, we'll cover both the graphical and console methods of logging in and starting the X Window System.

1.1 Log In, Please

Unlike some other operating systems, your Red Hat Linux system uses **accounts** to manage privileges, maintain security and more. Not all accounts are created equal: some have fewer rights to access files or services than others.

If you've already created a **user account**, you can skip ahead to Chapter 2, *GNOME in 60 Seconds (or so)*. If you created only the **root account**, read on to learn how to set up a user account.

Don't "root" around

Because your Red Hat Linux system creates one account (the root account) during installation, some new users are tempted to use only this account for all their activities. This is a bad idea; because the root account is allowed to do anything on the system, you can easily damage your system by mistakenly deleting or modifying sensitive system files. You may be tempted to forego creating and using a user account during or after installation, but you're playing with fire if you do.

1.1.1 Create a User Account

When you installed Red Hat Linux, you were asked to create a **root password** -- that is, a password for the root account, the **system administrator**. At that time, you were also able to create additional user accounts, which allow you to accomplish most tasks, without potentially harming your system -- as would be the case if you were to use your root account for everything.

If you didn't create a user account during the installation, that's your first task. Here's what you'll do:

- Log in from the console or from a graphical screen.
- Open a **terminal emulation window** (also called an **Xterm window** or **Xterm**) on the desktop.
- Create a new user account.
- Log out, then log in to the new account.

1.1.2 Logging In As Root

Regardless of whether you've chosen a graphical or console login screen, you'll have to supply a login account name and the password associated with that account.

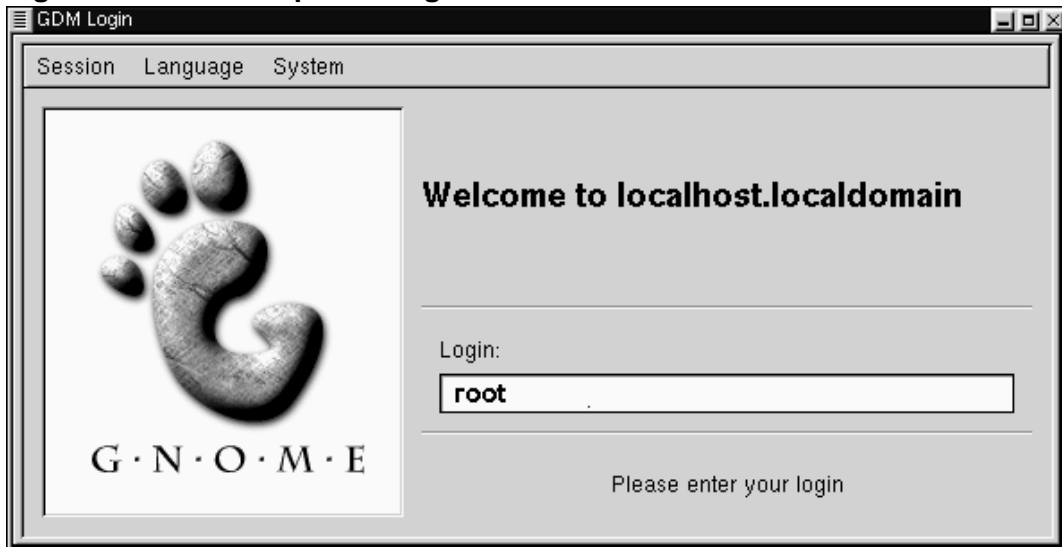
From a console screen, for example, you'll see something like:

```
Red Hat Linux release 7.0
Kernel 2.xx on an i686
localhost login:root
Password:yourrootpassword
```

Unless you've chosen to give your machine a **name**, as in a network setting, your machine will probably be called **localhost**.

To log in to the root account, at the login and password prompts, type **root** and the root password you chose when you installed Red Hat Linux.

Figure 1–1 The Graphical Login Screen



If you're using the graphical login screen, similar to Figure 1–1, *The Graphical Login Screen*, just type **root** in the box, press [Enter] and type in the password you created for the root account.

If you still see your console screen (instead of the graphical desktop) you can start the X Window System by typing **startx** as follows:

```
[root@localhost /root]#startx
```

To change your login screen

To find out how you can change from a console to a graphical login screen see Section 17.16, *Changing Login from Console to X at Startup*.

Once you start the X Window System, you'll find a desktop similar to Figure 1–2, *A GNOME Desktop* in GNOME or Figure 1–3, *A KDE Desktop* in KDE.

Figure 1–2 A GNOME Desktop

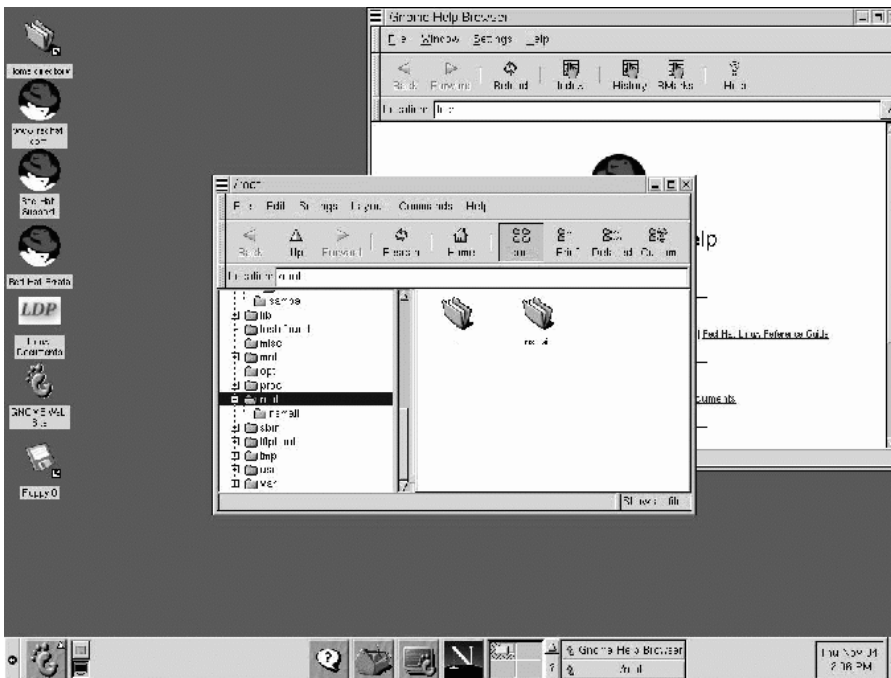



Figure 1–3 A KDE Desktop



1.1.3 Starting an Xterm

Both GNOME and KDE offer quick launch buttons on their **panels** to open an Xterm window.

Figure 1–4 The GNOME Panel



On the GNOME Panel, the button which launches an Xterm appears near the center

and looks like:



You can also find launchers to Xterms from the GNOME menu, under **Utilities**. Items which will open Xterms include **GNOME terminal**, **Regular XTerm**, and **Color XTerm**.

Figure 1–5 The KDE Panel

Similar to GNOME, the KDE Panel prominently features a quick launch button for

an Xterm. The launcher looks like: 

You can also find the launcher from the KDE main menu under **Utilities=>Konsole**.

Now, click on the **Xterm** button to open a window. You'll see the shell prompt inside the window, which will look like

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

Type `useradd`, then a space and the name of the new user account (name the account **newuser**, for example).

It may not appear that anything has happened, but you've just completed step one of creating the new account. All that remains is for you to give this account a **password**.

Picking account names

Often, user accounts are just variations on the user's name, such as `jsmith` for John Smith. You can choose whichever name you prefer, however, such as `musicman` or `ElivisK-ing`.

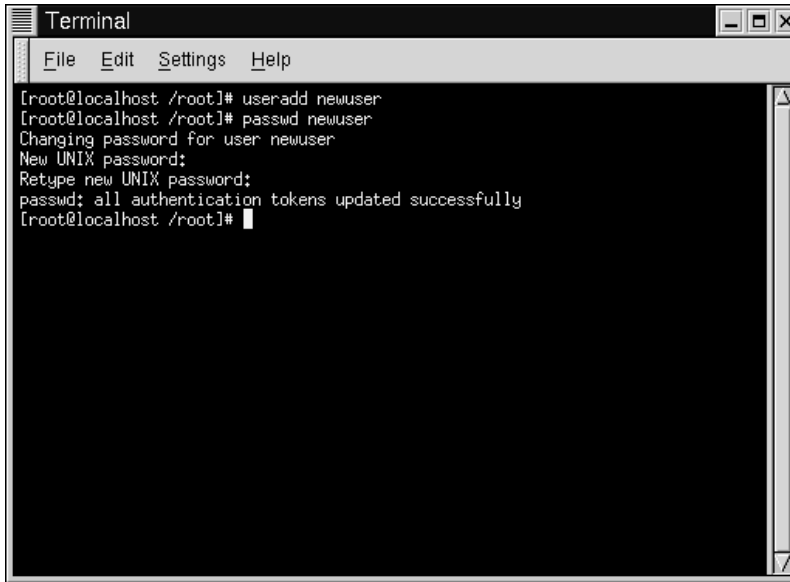
Now, on the next line, type `passwd`, a space and type the name of the new account for which you want to create the password (`passwd newuser`).

What is a secure password?

You can be fancy or plain when you pick a user account name, but take precautions when you choose a password. The password is the key to your account, so it should be both unique and easy for you to remember. Your password should be at least six characters (actually, it can be 256 characters long if you enabled MD5 passwords during the installation, though you probably don't need that many). You can mix upper- and lowercase letters, as well as numbers and characters. Avoid easy selections, such as `qwerty` or `password`. If you want to pick an easy-to-remember but somewhat unique password, consider a variation of a word, such as `a!rP18nE` for airplane. If you need more information about passwords, see Chapter 14, *Shell Prompt Basics*.

Next, you'll be asked to enter the password for the new user's account -- the prompt will state: `New UNIX password`. Enter a password that will be easy to remember and secure.

Type the password again for confirmation, and you'll see a message stating `passwd: all authentication tokens updated successfully`, which means that you've successfully created the new account.

Figure 1–6 Adding a User in an XtermA screenshot of an Xterm terminal window. The window title is "Terminal" and it has a menu bar with "File", "Edit", "Settings", and "Help". The terminal content shows the following commands and prompts:

```
[root@localhost /root]# useradd newuser
[root@localhost /root]# passwd newuser
Changing password for user newuser
New UNIX password:
Retype new UNIX password:
passwd: all authentication tokens updated successfully
[root@localhost /root]#
```


In Figure 1–6, *Adding a User in an Xterm*, you can see a sample of an Xterm window with the commands and prompts.

You can exit from the Xterm by clicking the **X** button on the upper right of the window, or by typing `exit` at the prompt.

1.1.4 Logging Out of Root

Now, you should log out of the root account and log in to your user account.

To log out from GNOME, click once on the **Main Menu Button** on the Panel and **drag** your mouse cursor to the first item, labeled **Log out** (as in Figure 1–7, *The Log out Selection*).

Figure 1–7 The Log out Selection

When the confirmation dialog appears (see Figure 1–8, *Logout Confirmation*), select the **Logout** option and click the **Yes** button. If you want to save the configuration of your Panel, as well as any programs which are running, check the **Save current setup** option, as well.


Figure 1–8 Logout Confirmation

Similarly, in KDE, you can log out from the menu button on the Panel (see Figure 1–9, *The Logout Entry in KDE*).

By default, the Panel in KDE also contains a quick launch button to log out; it's



located near the Taskbar, at the center of the Panel, and looks like:

Figure 1–9 The Logout Entry in KDE

You'll either return to the graphical login screen or the shell prompt, depending on how you chose to log in.

If you're returned to a non-graphical console prompt, just type `exit` at the prompt, as in:

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

Another way to exit

You can exit at the shell prompt either by typing the word `exit` or by using the key combination of `[Ctrl]-[D]`.

Now, you can log into your user account the same way you logged in as root.

2 GNOME in 60 Seconds (or so)

GNOME is the default desktop environment in Red Hat Linux 7.0. Like KDE, GNOME is flexible and user-friendly, providing both the power of your new Red Hat Linux system and a comfortable graphical user interface -- or **GUI**.

This chapter will help you get acquainted with some of the fundamentals of GNOME. You'll find greater detail about specific tasks -- such as connecting a printer and setting up your Internet account -- in later chapters.

You'll find the latest GNOME-related information at the the official GNOME website: <http://www.gnome.org>.

For a similar orientation to KDE, turn to Chapter 3, *KDE in 60 Seconds (or so)*.

If you want to use both

If you've got both GNOME and KDE on your system, you can often use applications from the other environment in your chosen environment. For example, you can use the KDE e-mail client, KMail, even though your chosen environment is GNOME. The same can be said for other applications which aren't strictly GNOME or KDE applications, such as those you might find online.

2.1 An Overview

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

- The Desktop -- The main workspace of your environment which incorporates everything, including the items which follow in this list (see Figure 2-1, *A GNOME Desktop* for an example). 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. See Section 2.2, *The Desktop*.
-

- **Folders** -- When you first log in, folders contain 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 2–2, *Folders and Launchers on the Desktop* for an example of folders and launchers on the desktop. See Section 2.3, *Folders*.
- **The Panel** -- The long bar across the bottom of the Desktop (for example, Figure 2–4, *The Panel at the Bottom of the Desktop*). The **Panel** serves as a quick launcher for applications and **applets** (mini-applications), as well as minimized applications you may already be running. See Section 2.4, *The Panel*.
- **Main Menu Button** -- Part of the **Panel**, this button, which looks like the stylized GNOME footprint (see Figure 2–7, *The Expanded Main Menu* for an example), holds launchers and submenus to all of the included applications and applets for your Red Hat Linux system. Some applications which you install will also be incorporated into this menu. You can add launchers to these menus, as well. See Section 2.5, *Main Menu Button*.
- **Help Button** -- By default, you'll find a button on the **Panel** which has a question mark on it. This button is a launcher to the GNOME Help Browser. See Section 2.6, *Finding Help*.
- **Multiple desktops** -- In the center of the **Panel** you'll find the Desk Guide, which shows you the number of **desktops** available to you. You can spread your work throughout multiple desktops, rather than crowd one desktop area with multiple applications. See Section 2.7, *Multiple Desktops*.
- **Switching environments** -- Although it's not something you'll find on your desktop, there's a utility which allows you to easily and quickly change your choice of environments. Called **Switchdesk**, you can use this utility to quickly and easily change from GNOME to KDE or vice-versa. See Section 2.8, *Switching Environments*.

2.1.1 Using Your Mouse

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.

Make your two-button mouse act like it has three buttons

A three-button mouse offers the greatest flexibility in GNOME. If you have a two-button mouse, however, and you selected three-button **emulation** when you installed Red Hat Linux, simply hold down both left and right mouse buttons at the same time to get the effect of a middle mouse button.

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

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, customizing the background 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.


2.1.2 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, and move or copy information and applications.

2.2 The Desktop

Your first view of GNOME will look something like Figure 2–1, *A GNOME Desktop*.


Figure 2–1 A GNOME Desktop



The desktop, which is also sometimes referred to as the **workspace**, is the area on which you run your applications and perform all your work. On your desktop are **launchers** and **folders** (see Figure 2–2, *Folders and Launchers on the Desktop*, for example).

- Launchers -- Icons which appear on the desktop. Click on the icon to start the application it's associated with. For example, clicking on a Netscape Navigator icon will launch the browser.
- Folders -- Containing contents of directories, launchers to documentation or applications.
- The folder **Home directory** -- Leads to your specific directory in `/home` on the system (for example, if your login directory was `newuser`, the **Home directory** folder on the desktop would lead to your login directory, `/home/newuser`) See Figure 2–3, *Contents of an Open Folder*.

Figure 2–2 Folders and Launchers on the Desktop



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

Launchers can also be shortcuts to websites. By double-clicking on the launcher labeled **Red Hat Support** on your desktop, for example, your Netscape Navigator browser will open and go to Red Hat's online support site.

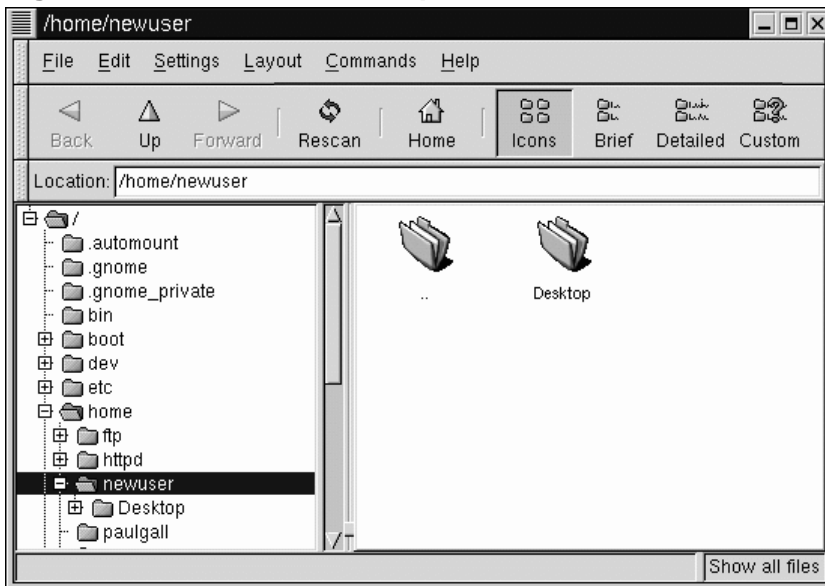
2.3 Folders

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 display its contents.

Folders correspond to directories on your system. Your default desktop includes a folder showing you the contents of your /home directory (see Figure 2–3, *Contents of an Open Folder*). Although you can specify different ways to view the contents of

your /home and other directories on your system, in general, the left pane of the file manager shows you the directory **tree** while the right pane shows you the contents of that directory.

Figure 2–3 Contents of an Open Folder



You can start the file manager by double-clicking on a folder on your desktop or from the main menu, under **File Manager**.

What is the file manager?

The file manager is actually an application called the GNU Midnight Commander, or GMC. GMC is GNOME's version of the popular Midnight Commander file manager. You can learn more about the File Manager in Section 10.1, *The GNOME File Manager* or from the *GNOME User's Guide* at the GNOME website.

2.4 The Panel

Like much of GNOME, the **Panel** is highly configurable (see Figure 2-4, *The Panel at the Bottom of the Desktop*).

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 GNOME Pager and clock applet.

Like KDE, you can **hide** the **Panel** by clicking on one of the arrows at either of its borders. You can also **auto-hide** the **Panel**, making it appear to sink beneath the frame of your screen (it will return when you hold your cursor over the minimized space). To enable the auto-hide feature, right-click on the **Panel**, select **panel, properties, hiding policy** and choose **auto hide**.


Elements of the GNOME Pager

The GNOME Pager consists of the Desk Guide, which is a visual representation of your desktops, and the Tasklist, which shows your running applications -- either minimized or opened on the desktop.

For more on the Desk Guide and GNOME Pager, see Section 2.7, *Multiple Desktops*.

Other applets which can enhance productivity and enjoyment range from a mail checker and a CD player to system monitors. Applications, folders and applets can be docked on the **Panel**, so that they can be launched with a single click of the button.


Figure 2–4 The Panel at the Bottom of the Desktop



To add an applet:

- Right-click on the **Panel**; select **Applets** (see Figure 2–5, *Adding an Applet*)

Figure 2–5 Adding an Applet

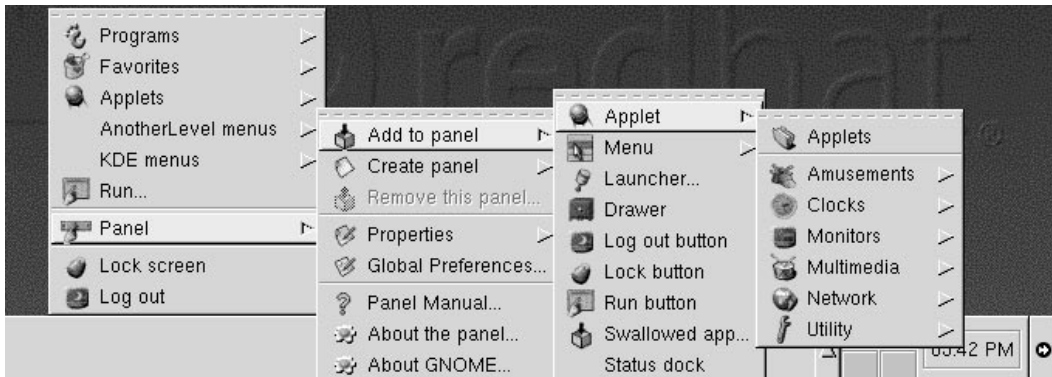


*Select the type of applet
you want to add*

Right-click on the Panel

- Right-click on the **Panel**; select **Panel => Add to panel** (see Figure 2–6, *Adding an Applet*)

Figure 2–6 Adding an Applet



- Choose **Applets** from the main GNOME footprint menu


You can also customize the **Panel**, add more than one **Panel** to your desktop, or convert the current orientation of the **Panel** from horizontal to vertical. By clicking on the arrows which appear at either end, you can "shrink" the **Panel** to appear in a corner of the Desktop, rather than extend across the screen (as in Figure 2–4, *The Panel at the Bottom of the Desktop*).

Switch tasks quickly

One way to bring currently running applications to the front is to press [Alt]-[Tab]. To scroll through the tasks, hold down the [Alt] key, while pressing the [Tab] key in succession. When you've found the task you want to bring to the front, stop pressing the [Tab] key and release the [Alt] key.

2.5 Main Menu Button

The stylized GNOME footprint on the **Panel** is called the **Main Menu Button**. Left-click on the button, and you'll see a menu similar to Figure 2–7, *The Expanded Main Menu*.

Figure 2–7 The Expanded Main Menu

From here, you can start all the included applications and applets for GNOME; you also have access to KDE applications, under **KDE menus**, and to applications you would find in the window manager AnotherLevel (provided that these applications were installed in the first place).

These submenus give you access to a full range of applications on your system -- in GNOME, KDE and AnotherLevel. If you prefer to use the GNOME environment, but would rather use some non-GNOME applications, you can easily choose them from the **Main Menu Button**.

From the **Main Menu Button**, you can also log out, run applications from a command line and **lock your screen**, which runs a password protected screensaver.

2.6 Finding Help

The GNOME Help Browser (see Figure 2–8, *The GNOME Help Browser*) is one of the more useful utilities at your command.

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 learning 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 system. These help documents are covered more extensively in Chapter 14, *Shell Prompt Basics*.)

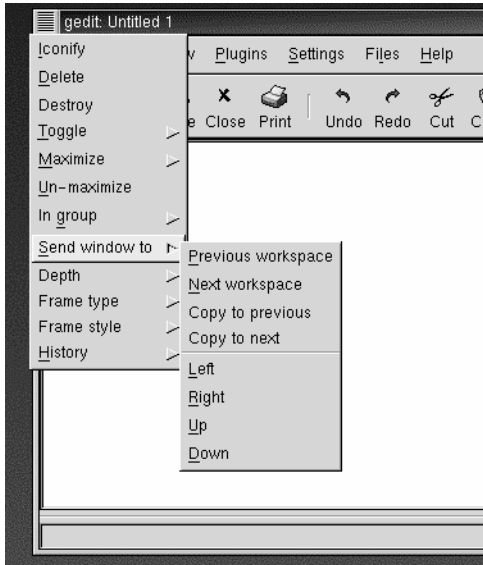
Figure 2–8 The GNOME Help Browser



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** on the **Main Menu Button**.

2.7 Multiple Desktops

Rather than crowd your work on one screen, you can use **virtual desktops**, which extend your workspace to multiple desktop areas. Say you have your browser open on one desktop; if you want to use a word processor and e-mail program without crowding everything together, you can have each application run on its own desktop.

Figure 2–9 Moving an Application to Another Desktop

For example, using GNOME and the **sawfish window manager** (which literally *manages* how your windows are displayed), you can move a window from one desktop to another by clicking on the application's title bar and "dragging" the application beyond the left, right, top or bottom border of your screen.

By default, four desktops are available to you. However, using the configuration utility for your window manager, you can modify that number. To run Enlightenment's configuration program, middle-click on your desktop (or click your right and left mouse buttons simultaneously), and select **Customize**. Go to the entry called **Workspaces** to increase or decrease the number of available desktops.

Figure 2–10 The GNOME Desktop Guide

The Desk Guide (see Figure 2–10, *The GNOME Desktop Guide*) is a GNOME utility which allows you to switch desktops (or workspaces) by clicking on the corresponding desktop.

You can use the GNOME Desktop Guide, the collection of "panes" on the **Panel**, to see where running applications are located on your desktops -- they will appear as small, raised squares inside a desktop pane.

Part of the Pager


The Desktop Guide is one part of the GNOME Pager; the other part of the Pager is the Tasklist. You can add new Tasklists, Pagers, or Desk Guides by right-clicking on the **Panel** and moving your cursor to **Add applet => Utility** and selecting the applets from the menu.

2.8 Switching Environments

In addition to changing your graphical environment at the login screen, you can "switch" to another environment when you're already logged in.

You can change environments by using the **Switchdesk** utility (see Figure 2–11, *The Switchdesk Utility*). Whether you're using GNOME, KDE or AnotherLevel, you can select your new environment in **Switchdesk** and log out. When you restart your session, you'll find your new environment.

In GNOME, one way to find **Switchdesk** is from the **Main Menu Button** (go to **Programs => System => Desktop Switching Tool**).

Figure 2–11 The Switchdesk Utility

Other ways to start the utility include:

- Type **switchdesk** at the shell prompt.
- Open the **Run program** item from the **Main Menu Button** and type **switchdesk** in the window.

Your changes won't be reflected immediately. Instead, you'll have to log out of your current X session. The next time you start X, you'll work in your new environment.

3 KDE in 60 Seconds (or so)

Like GNOME, KDE is an appealing and easy to configure graphical environment.

Formally called the K Desktop Environment (the K doesn't represent any word, according to the website FAQ), the goals of the KDE Project are to lessen the Linux/UNIX learning curve for many users, and to present an appealing environment for everyone. Since its foundation in October 1996 by Matthias Ettrich, KDE has garnered praise from new and seasoned Linux users alike.

Hundreds of free software developers from all over the world have worked on the KDE Project, and continue to refine and enhance the environment.

Like the chapter about the basics of GNOME (Chapter 2, *GNOME in 60 Seconds (or so)*), this chapter will provide a quick overview of KDE fundamentals. Details, such as configuring your desktop, working with the file manager and using specific KDE tools, will be covered later in this book.

For the latest documentation on KDE, visit the project's website: <http://www.kde.org>.

If you want to use both

If you've got both GNOME and KDE on your system, you can often use applications from the other environment in your chosen environment. For example, you can use the Gnumeric spreadsheet application in KDE, even though the application is part of the bundled GNOME applications.

3.1 An Overview

When you first start KDE, among the first items you'll see are:

- The Desktop -- The main workspace of your environment (see Figure 3–1, *A Typical KDE Desktop* for an example); the space on which you place windows for
-

running applications, icons for starting programs, folders for programs or documents. See Section 3.2, *The Desktop*.

- The **Floppy**, **CD-ROM**, **Printer**, and **Trash** icons -- These are links (or **shortcuts**) which allow you easy access to your floppy and CD-ROM drives or let you "throw away" items by dropping them into the trashcan. See Section 3.2.1, *The Trash, Printer, and Drives*.
- Folders -- In the folder labeled **Templates**, you'll find links which you can **drag** to your desktop to create new desktop items. The folder called **Autostart** is where you can place applications which you want to start automatically whenever you log in (see Section 3.3, *Folders*).
- The **Panel** -- The long bar across the bottom of the Desktop (for example, in Figure 3-1, *A Typical KDE Desktop*). The **Panel** serves as a quick launcher for applications and **applets** (mini-applications). See Section 3.4, *The Panel*.
- **Multiple desktops** -- In the center of the **Panel** you'll find the **Switch Desktop** buttons, which show you the number of **workspaces** available to you. You can spread your work throughout multiple desktops, rather than crowd one desktop area with multiple applications. See Section 3.4.1, *Multiple Desktops*.
- **Main Menu K Icon** -- Part of the **Panel**, at the left (see Figure 3-9, *The Main Menu K Icon with Expanded Menu* for an example), leads to launchers and submenus for included applications and applets on your Red Hat Linux system. Many applications which you install, such as **StarOffice**, will also be incorporated into this menu. See Section 3.5, *The Main Menu K Icon*.
- The **Taskbar** -- Appearing at the top of the desktop by default (though you can change this), the **Taskbar** is a thin bar which shows your currently opened windows -- both minimized and on the desktop. See Section 3.6, *The Taskbar*.
- Finding Help -- KDE includes a detailed set of help files for navigating and customizing the environment. See Section 3.7, *Finding Help*.
- Switching Environments -- Although it's not something you'll find on your desktop, there's a utility which allows you to easily and quickly change your choice of environments. Called **Switchdesk**, you can use this utility to quickly and easily

change from GNOME to KDE or vice-versa. See Section 3.8, *Switching Environments*.

3.1.1 Using Your Mouse

Click once

Unlike GNOME, you start applications and open folders by clicking just once on the associated icon.

Using the left mouse button selects an item to open or drag to another location. To drag the item, hold down the mouse button, and continue holding it down while dragging to the new location. Once at the desired location, release the mouse button to "drop" the item.

You can also select multiple items by holding down the cursor and dragging it over the items. Once you release the mouse button, each of the selected items will be highlighted.

If there are multiple items you want to move or copy, but they're not easily aligned for quick selection, you can highlight each item by holding down the [Ctrl] key while clicking once on each icon.

When you **drop** the files on their desired location -- such as a folder -- a menu will appear asking whether you want to **Copy**, **Move** or **Link** the files at the selected location.

Make your two-button mouse act like it has three buttons

As with GNOME, a three-button mouse offers the greatest ease of access to menus and navigation in KDE. 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 to emulate the "middle" button of a mouse.

Pressing the middle button of your mouse on the desktop will allow you to switch desktops by selecting desktops one through four (you can rename the desktops, as you'll learn in later in this book).

Using the right mouse button reveals a brief menu of actions which you can take. Right-clicking on the desktop shows you options such as adding desktop links -- which start applications, visit websites or open documents. You can also rescan or unclutter the desktop.

Also, among the options when you right-click on the desktop, you'll find a link to helpful documentation on using the desktop itself.

Right-clicking on a folder or application reveals a menu related to actions you can take with your selection, such as moving the item to the **Trash**, revising the item's properties and more.

3.1.2 Drag and Drop

You can **drag and drop** items in KDE, 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.


When you drag an item by holding down the left mouse button, then release the button over the desired location, you'll find a menu of actions you can take. You can copy the item into the location (which will preserve the original item) you can move the item to its new location, or you can create a link from the item in its new location to the original item. This last option is somewhat similar to copying; however, when you copy the item, the original item remains unchanged if you modify the copy. When you link an item, any changes you make to the item in one location are reflected in the item at its original location.

You can also drag applications, such as **Netscape Navigator**, from the desktop and drop them onto the **Panel** to use as a quick launcher. (To remove an item from the **Panel**, right-click on the item, and select **Remove**. Removing an application from the **Panel** will only remove the link and icon, not the actual application.)

3.2 The Desktop

When you first start KDE, your desktop will look something like Figure 3–1, *A Typical KDE Desktop*.

Figure 3–1 A Typical KDE Desktop



The desktop is your main work area. It is the area on which all of your running applications will be placed -- throughout one or multiple desktops.

Like GNOME, the KDE desktop can be easily configured.

To start or open any of the applications or documents, *single-click* on an icon on the desktop. The associated program or document will open after the icon is briefly highlighted.

Dressing it up

Like so much of KDE, you can customize the look of your desktop, either through the use of preferred colors or pictures, or through the use of **themes**. To learn more about themes in KDE, visit the KDE website (<http://www.kde.org>).

3.2.1 The Trash, Printer, and Drives

Among the first items you'll discover on your desktop will be the icons representing your printer, CD-ROM and floppy drives, as well as a **Trash** can (see Figure 3–2, *Drives, Trash and Printer Icons*).

Figure 3–2 Drives, Trash and Printer Icons



Both the icons for "floppy" and "cdrom" provide shortcuts to those devices. You can access either of these devices by left-clicking once on one of the icons.

Be aware of formatting

If you place a floppy diskette in the drive, then select the **floppy** drive on your desktop, you may receive an error message such as "**wrong fs type, bad option, bad superblock on /dev/fd0...**"

Chances are good that the floppy you've placed in the drive is not formatted with the ext2 filesystem -- the Linux filesystem format -- but rather another filesystem format, such as MS-DOS. For more information about how to access an MS-DOS formatted diskette, as well as an ext2 filesystem, see Section 17.2, *Using a Diskette*.

When you right-click on the drive icons, you will see several options for working with these drives, such as **Delete** and **Move to Trash**, which will remove the icon from your desktop (and which you probably won't want to do); **Copy**, which allows you to copy the icon to another location; **Open with**, which you should probably avoid for now; and **Mount**.

"Mounting" a device

Ordinarily, you'll be able to use your drives just by clicking on them. In UNIX and Linux, however, that's not always the case: you've often got to **mount** the drives. When you mount a floppy or CD-ROM, you're making that device's contents available to you. Note, however, that whenever you mount a drive, you must also **unmount** the drive when you're finished.


The **Trash** is where you can drag and drop your unwanted items. Right-click on the trashcan and select **Empty Trash Bin** to delete the items from your system.

The **Printer** shows you the **print queue**, meaning the jobs which are printing or scheduled to print. This is a handy tool for monitoring the progress of your printing tasks.

3.3 Folders

When you first start KDE, you'll see "file folders" on your desktop. Among the more important folders are **Templates** and **Autostart** (see Figure 3–3, *File Folders on the Desktop*).

Figure 3–3 File Folders on the Desktop




To open a folder, single-click on it with your left mouse button.

In the **Autostart** folder, you can drag and drop applications which you would like to open or start automatically when you begin your KDE session.

In the **Templates** folder, you'll find basic file-types which you can use to drag onto your desktop or into other folders, then customize to suit your needs (see Figure 3–4, *A View of the Templates Folder*).

Figure 3–4 A View of the Templates Folder



Among the templates you can choose are:

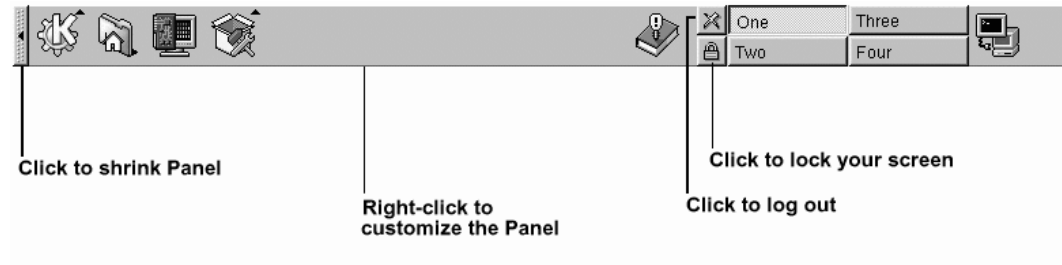
- **Ftpurl** -- Creates a link to **FTP** sites. (FTP stands for **File Transfer Protocol**, a means of exchanging files, such as on the Internet.) Single-click on the link to access the site.
- **Mime Type** -- Creates a link to an application or text file. (MIME means Multi-purpose Internet Mail Extensions; they're used to describe different data so that the operating system knows how to work with it when that data is selected.) You won't require this template often, since many MIME types are already available with KDE.
- **Program** -- Creates a link to start a program.
- **URL** -- Creates a link to a favorite website, like a detachable bookmark.

To use these templates, left-click on your chosen template, drag it to your desktop or into another folder, then customize it with the appropriate name and information.

3.4 The Panel

The **Panel** stretches across the bottom of the desktop.

Figure 3–5 The Panel



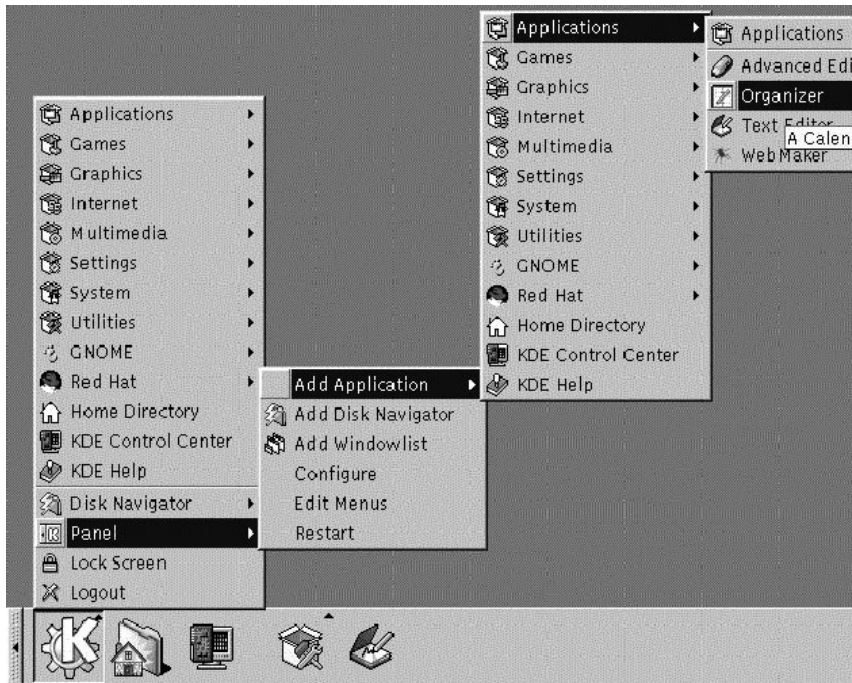
The **Panel** is very configurable. You can add and remove applications and applets, which will allow you to access them easily.

As in GNOME, you can **hide** the **Panel** by choosing to **auto-hide** it, making it appear to sink beneath the lower frame of your screen (it will return when you hold the cursor over the minimized space).

You can also hide the **Panel** by clicking on one of the arrows on either of its borders; depending on the arrow you select, that will be the direction in which the **Panel** will shrink to hide.

Applications and utilities can easily be added to the **Panel** (see Figure 3–6, *Adding an Application to the Panel*).

Figure 3–6 Adding an Application to the Panel



To add an application:

- Click on the **Main Menu K** icon and highlight the **Panel** entry in the menu.
- In the submenu which opens, highlight **Add Application**.
- Highlight one of the submenus which open and select the corresponding application, game or utility you wish by dragging your cursor to the item until it's highlighted.
- Once your choice is highlighted, release the mouse button. The item will be added to the **Panel**.

3.4.1 Multiple Desktops

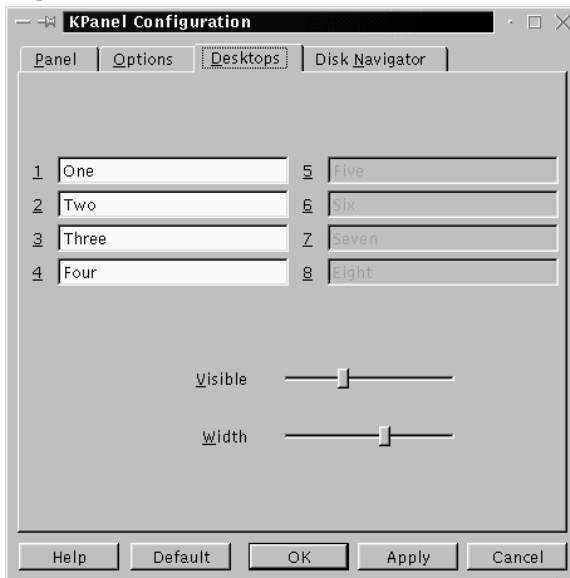
By default, KDE offers you four different desktops over which you can spread your work. Each desktop can hold icons, open applications and individually customized backgrounds.

For example, while you're writing e-mail on the first desktop, you can have Netscape Navigator running on desktop two, a word processor open on desktop three, and so on.

You can change the number -- and names -- of desktops available under KDE. To make these adjustments:

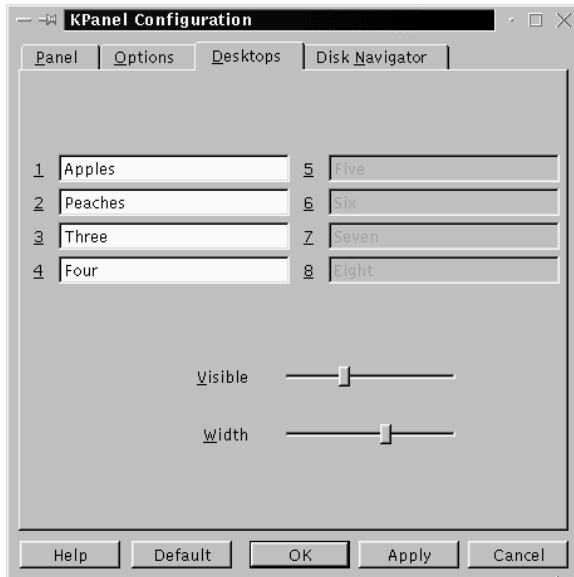
- Right-click on the **Panel**; you'll see a limited menu of actions you can take.
- Select **Configure**; the KDE **Panel** configuration tool will open.
- Click on the tabbed entry marked **Desktops** (see Figure 3-7, *The Desktops Tab in the Panel Configuration*).

Figure 3-7 The Desktops Tab in the Panel Configuration



To change the names of your desktops, from **One**, **Two**, and so on, replace the desktop "names" by typing in your preferred names in the appropriate boxes (see Figure 3–8, *Renaming Desktops*).

Figure 3–8 Renaming Desktops



You can increase (or decrease) the number of desktops available to you by adjusting the **Visible** slider of the **Desktops** dialog. For more desktops, drag the bar to the right; for fewer desktops, drag the bar to the left.

In similar fashion, you can adjust the width of the desktop buttons on the **Panel**.

Switch desktops quickly

Want to change desktops without the hassle of pointing and clicking? Just use the keyboard combination of the [Ctrl] and Function keys. For example, [Ctrl]-[F2] will switch you from desktop one to desktop two, [Ctrl]-[F3] will take you to desktop three. (If you rename your desktops, remember which number represents the associated name!)

3.5 The Main Menu K Icon

Similar to the **Main Menu** of GNOME (and many other graphical environments), the main menu in KDE, at the **K** icon, incorporates all of the available KDE applets and applications, as well as many other applications on your Red Hat Linux system.

Figure 3–9 The Main Menu K Icon with Expanded Menu



You can find the icon on the left of the **Panel**. You open the main menu by clicking once on the icon. As in Figure 3–6, *Adding an Application to the Panel*, you'll see the top-most level of menu headings first. These headings are the major groups into which the applications are placed. For example, to open **KOrganizer**, the KDE calendar and appointment book, you would open **Main Menu K**, drag your mouse to highlight the **Applications** menu listing, then highlight and click once on the **Organizer** entry in the submenu which opens.

Also available from the main menu, you'll find:

- The **KDE Control Center**, an invaluable resource for customizing your environment and locating system information.
- A **Home Directory** listing which, when selected, will open a browser showing you all the files in your account.
- A menu called **Disk Navigator**, which provides quick access to recent applications as well as important directories in the filesystem.
- The **Lock Screen** and **Logout** icons, which will password protect your desktop and allow you to log out of your account.

3.6 The Taskbar

The **Taskbar** displays all running applications -- both minimized and on the desktop (see Figure 3–10, *Applications on the Taskbar*).

Figure 3–10 Applications on the Taskbar



Like GNOME, Windows 9x and other graphical environments, the **Taskbar** is an easy way to keep track of and access running applications.

You can keep the **Taskbar** running in a separate location from the **Panel** or you can have them both located in the same place on the desktop. The default configuration places the **Taskbar** at the top of the desktop while the **Panel** runs at the bottom of the desktop.

You can maximize running applications or bring them to the front of your working windows by clicking on the associated item on the **Taskbar**.

Switch tasks quickly

Another way to bring minimized or background windows to the front is to use the [Alt] and [Tab] keys. To pick an item from the **Taskbar**, hold down [Alt]-[Tab]. To scroll through the tasks, hold down the [Alt] key, while pressing the [Tab] key in succession. When you've found the task you want to maximize and bring to the front, stop pressing the [Tab] key and release the [Alt] key. This procedure works only on your current desktop, so if you've got applications running on other desktops, you'll have to switch to that desktop in order to maximize it with these keys.


3.7 Finding Help

KDE Help is the online help browser for getting the most out of KDE, as well as for many features of your Red Hat Linux system.

Figure 3–11 The Opening Screen of the Help Browser



There are several ways to access KDE Help, but the easiest are:

- From the **Panel** -- A single click on the KDE Help icon, which looks like: 
- From the **Main Menu** -- Listed as **KDE Help** next to the help icon.
- By right-clicking on the desktop and selecting the **Help on desktop** entry in the brief menu.

The opening screen of the KDE Help browser will appear like Figure 3–11, *The Opening Screen of the Help Browser*. From this main page, you can view help documentation on using and configuring the desktop, working with KFM, the KDE file manager, and more.

The main contents of the help browser are:

- *Welcome to the K Desktop Environment* -- Introductory material about the KDE Project.
- *KDE Help Contents* -- Help documentation such as **man pages** and **info pages**, as well as specific help on KDE applications.
- *Getting the most out of KDE* -- Quick-start tutorials, information on KFM and kwm, the **window manager** for KDE (a window manager does just that: manages window display and behavior in an environment).

A useful extension to the help browser is the ability to search through the help documentation installed on your system. You can search through KDE documentation and through man pages (eventually, you'll also be able to search through info pages, although this feature is not yet available).

You'll find a link to the search tool in the **KDE Help Contents** section of the help browser. To open the search tool, select the link **Search for Keyword**.

Figure 3–12 The KDE Help Search Tool



In the search tool (see Figure 3–12, *The KDE Help Search Tool* for an example screen), type in the keyword, subject or significant phrase for which you want help. Select from KDE-specific documentation (**KDE Documentation**) or man pages (**Online Manuals**), which are system-related help documentation, or both categories. Click on **Submit Search** or press [Enter] to begin searching.

When you find useful material, you can print the documentation by going to **File => Print** on the menu, or by pressing [Ctrl]-[P].

3.8 Switching Environments

In addition to changing your graphical environment at the login screen, you can "switch" to another environment when you're already logged in.

You can change environments using the **Switchdesk** utility (see Figure 3–13, *The Switchdesk Utility*). Whether you're using GNOME, KDE or AnotherLevel, you can select your new environment in **Switchdesk**, then log out of your session. When you restart your session, you'll find your new environment.

Figure 3–13 The Switchdesk Utility



Start **Switchdesk** from the shell prompt in an Xterm window. You'll find a quick



launcher to an Xterm on the **Panel**; it looks like:

To use **Switchdesk**, perform the following tasks:

- Open an Xterm window.
- Type **switchdesk** at the prompt.
- Select your new environment and click **OK**.
- Close the Xterm window by typing `exit`.
- Log out and log in again.

The next time you start your X session, you'll be in your newly selected environment.

4 Logging Out and Shutting Down

By now, most computer users know that shutting down their system involves a little bit more than just tapping a power switch. The same is true for your Red Hat Linux system.

But just because you're done working doesn't mean you must shut down. For example, if you're sharing your computer with another user, you can leave your machine running and simply **log out** of your account, which can avoid a time-consuming reboot.

There are similarities and differences to the procedure of logging out and shutting down in both GNOME and KDE. Let's cover the basics first in each environment.

4.1 Quitting from GNOME

When you log out, GNOME allows you to simply log out (and leave the system running), restart, or halt the system completely.

Figure 4–1 The GNOME Log Out Confirmation

To log out from the **Main Menu** on the **Panel**, click on the **Main Menu Button** and select **Log out**.

In the confirmation dialog which opens (see Figure 4–1, *The GNOME Log Out Confirmation*), you can choose to:

- **Logout** -- logs you out of your account and returns you to the login screen, leaving the system running.
- **Halt** -- logs you out of your account and shuts down the system.
- **Reboot** -- logs you out of your account and restarts the system.

Select any of the options and press the **Yes** button to continue. Also, select the **Save current setup** option if you want to **save your session**. Saving your session will preserve your current configuration of the **Panel** and save the programs you might have open.

If you don't wish to proceed, choose the **No** button to continue with your GNOME session.

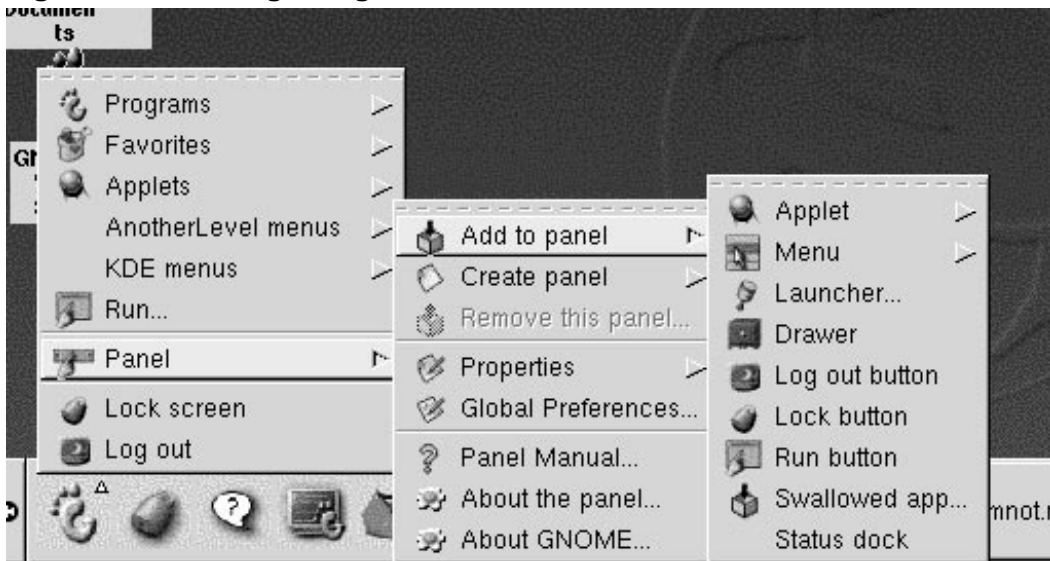
Typing your password

When you select either **Halt** or **Reboot**, you'll be asked for your account's password before proceeding. In the space provided, just type in the password you used when you logged in for your current session.

4.1.1 Adding a Log Out Button

You can add a quick launch **Log out** button to the **Panel**.

Figure 4–2 Adding a Log Out Button



To add the button to your **Panel**, go to the **Main Menu** => **Panel** => **Add to panel** and click on **Log out button** (see Figure 4–2, *Adding a Log Out Button*). Then, just click on the launcher on the **Panel** to log out.

4.1.2 Shutting Down from the Run program menu


An additional way to shut down or restart your machine is to select the **Run program** item from the **Main Menu** and, in the dialog which opens, type `shutdown -r now` or `shutdown -h now`. The `-r` means reboot, and will restart your machine; the `-h` means halt, and will shut down the system.

We'll discuss these two commands a little later in this chapter.

4.2 Quitting from KDE

There are several easy ways to log out from your KDE session:

Figure 4-3 Logout from the Main Menu



- From the **Main Menu** -- Select the **Logout** item (see Figure 4-3, *Logout from the Main Menu*).

- From the **Panel** -- Select the **Logout** icon, which looks like:



- From the desktop -- Right-click on the desktop and select **Logout** from the menu.

Like GNOME, KDE offers you the chance to save your current settings, which will preserve your panel settings and keep your currently open applications.

If you're working in an application and you haven't saved your work when you log out, a dialog will inform you that you'll lose your unsaved material when you log out.

When you see this dialog, you can simply select the **Cancel** button, save your work, then log out again. If you continue without saving your work, you'll lose the data.

4.3 Shutting Down at the Shell Prompt

Like the **Run program** option in GNOME, you can use the `shutdown` command to halt or restart your system from the shell prompt in an Xterm window. Whether you're working in GNOME, KDE or the **console mode** (non-graphical), you can shut down or restart your system from the shell prompt.

Figure 4-4 The Password Dialog



To halt or restart your system from the shell prompt:

- Open an Xterm window -- From the **Panel**, click on the **Terminal Emulation** or **Console** launcher.
- At the shell prompt, type `shutdown -r now` or `shutdown -h now` and enter your user account password in the dialog box which appears (see Figure 4-4, *The Password Dialog*). The `-h` means halt and will shut down the system; the `-r` means reboot and will restart the system.
- In console mode, type `shutdown -r now` or `shutdown -h now` and enter your user account password after the `Password` prompt.

Make certain that you've saved your work before halting or restarting your system from the shell prompt. Running applications will be closed and you won't have the option of saving your work or your session.

"Killing" an X session

If you're having difficulty quitting your X session -- for example, if you get no response when you click on the **Log out** button in KDE -- you can swiftly **kill** the X session and return to the login screen by pressing the [Ctrl]-[Alt]-[Backspace] keys. This procedure should be used only as a last resort, however. *Make sure you close your running applications and save your work before logging out this way!*

More about shutting down

You can specify the delay between when you issue the `shutdown` command and its execution. Replace the word `now` with a "+" and a number; for example `shutdown +h +5` will shutdown and halt the system after five minutes pass.

5 Setting Up Your Printer

Most new Red Hat Linux users want to immediately set up their printer.

Here, we'll cover some basic steps to configure a parallel printer on your system. For the most part, the process isn't too complicated, regardless of whether you're using GNOME or KDE.

Is it compatible?

Although many printers work well with Linux, not all printers are compatible. To make certain your printer is compatible with Linux, refer to *The Hardware Compatibility List* at <http://www.redhat.com/support/hardware/>.

Another useful site which features printer compatibility lists as well as useful hints and instructions is *The Linux Printing HOWTO*, at <http://www.linuxprinting.org/howto/>.

Because Linux is a multi-user, multi-tasking operating system, it can handle a variety of printing jobs at once. You could have multiple users on the system, all of whom want to send their work to a printer. Linux dutifully handles these requests by **spooling** them -- storing and printing the files in the order in which they were received.

Even if you're the only user on your system, you're still going to send your work to the **print spool**, where the files wait until it's their turn to be printed.

There are some terms and conditions you should know:

- **lpd** -- The **line printer spooler daemon**. Normally starts when the system starts. Manages the print jobs.
 - **lpr** -- Takes each print job and submits it to lpd.
 - **Input filters** -- Manage printer-specific formatting.
 - `/etc/printcap` -- This is the configuration file which defines your printer and other specifics.
-

- `/var/spool/lpd/` -- This is the location of the print spool, which is where files await printing.
- **Printtool** -- A graphical utility you will use to configure and modify `/etc/printcap`.

Read the "man page"

You can learn more about `lpd` and related commands and daemons by typing `man lpd` at a shell prompt. To move forward through the man page, use the [Space]; to move back, press the [B] key, and to quit, press the [Q] key.

Regardless of whether you're using GNOME or KDE, you can configure your printer with **Printtool**. **Printtool** provides a graphical means for setting up and managing your printer.

5.1 Using Printtool

The file `/etc/printcap` defines how `lpr` should process the print jobs. If you were to take a peek at the file, you'd find it looks something like the following excerpt:

```
##PRINTTOOL3## REMOTE POSTSCRIPT 600x600 letter {} PostScript Default 1
acct|xerox-ac:\
    :sd=/var/spool/lpd/acct:\
    :mx#0:\
    :rm=xerox-ac.redhat.com:\
    :rp=lp:\
    :if=/var/spool/lpd/filter.d/filter:
```

It's not pretty, and unless you know exactly what you're doing, you shouldn't edit the file. Instead, use the utility **Printtool** to set up the printer (and `/etc/printcap`) for you.

To use **Printtool**, you'll have to "be root."

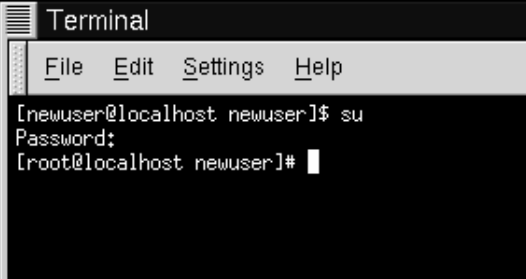
Whenever you make changes to your system configuration -- such as adding or modifying hardware settings, installing or uninstalling software -- you must be **superuser**, or root.

To become root from your user account:

- Open an Xterm window;
- Type `su` at the prompt;
- Type your root password. (You won't see it "echoed" on the screen when you type it.)

See Figure 5–1, *Switching to Root* for an example of how to become root.

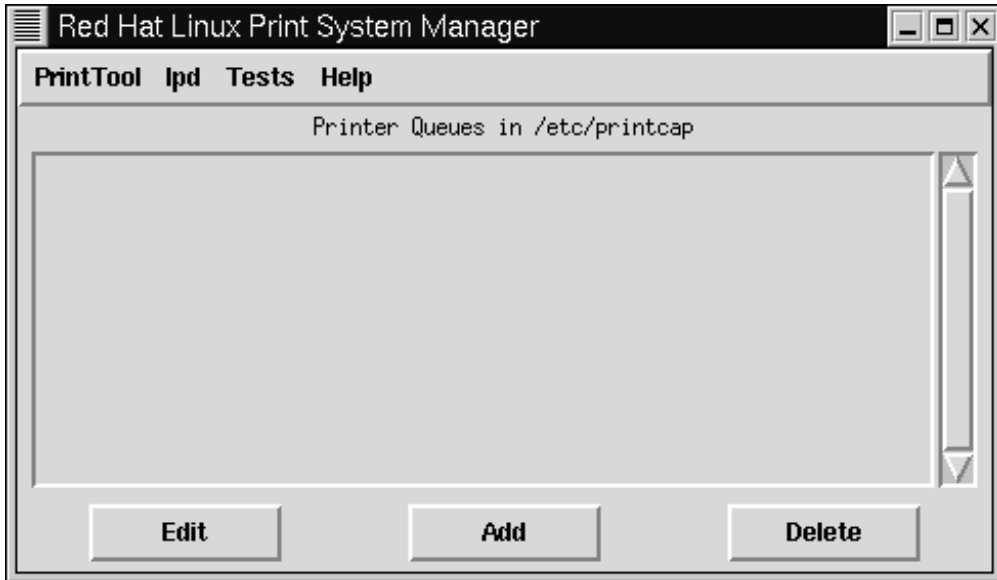
Figure 5–1 Switching to Root

A terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal text shows a user prompt "[newuser@localhost newuser]\$" followed by the command "su". A "Password:" prompt is shown, and the prompt then changes to "[root@localhost newuser]#" with a cursor.

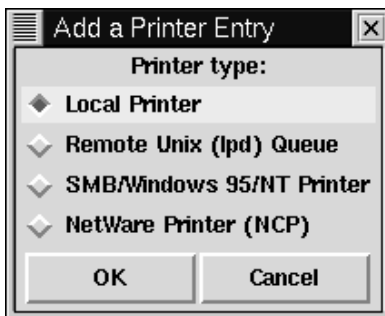
As you can see in the previous figure, when you're root your prompt changes from a `$` to a `#`. Type `printtool` -- one word, without capitalization.

When the utility starts, you'll see a window similar to Figure 5–2, *The Printtool Window*. At first, you won't find any printers listed, because you haven't added a printer to `/etc/printcap`.

Click on the **Add** button to add your printer.

Figure 5–2 The Printtool Window

Next, you'll be asked what kind of printer you want to add (see Figure 5–3, *Selecting the Type of Printer to Add*).

Figure 5–3 Selecting the Type of Printer to Add

Choose from the following:

- **Local Printer:** Choose this if your printer is directly connected to the machine on which you're working (called the **local machine**).
- **Remote Unix (lpd) Queue:** A printer attached to another machine on which an lpd daemon is running. Choose this if you want to print to a network printer on a LAN.
- **SMB/Windows 95/NT Printer:** A printer which is attached to a machine running Windows or Samba. (Samba is a suite of programs which allows Windows machines to share files and services with UNIX and Linux machines on a network.)
- **NetWare Printer (NCP):** A printer on a machine running Novell NetWare or compatible server software -- two networking applications for interacting with machines running Novell NetWare.

What is a daemon?

You may have seen that lpd is called the line printer spooler daemon. A daemon is a small program that runs in the background. Daemons usually perform critical system tasks, and usually start when the system starts. (Most daemons, also, have names that end in a d, such as crond, httpd, lpd, and so on.)

For this example, we'll add a local printer. Highlight the entry **Local Printer** and select the **OK** button.

When you select your printer type, Printtool will attempt to detect any printers which are attached to your parallel port; you'll then see the results (see Figure 5-4, *Detecting a Printer Port*). Click on the **OK** button to continue to the main Printtool window.

Figure 5–4 Detecting a Printer Port

The `/dev/lpx` Device Files

Linux refers to parallel ports differently than DOS. In DOS, a printer located at LPT1 is known in Linux as `/dev/lp0`; LPT2 in DOS is `/dev/lp1` in Linux, and so on.


In our example, a parallel port printer was detected at `/dev/lp0`. Click on the **OK** button to continue.

In the next dialog (see Figure 5–5, *Editing a Local Printer*), you can specify details about your printer, such as a name of the printer and its related spool directory, the location of the printer device and other options.

In the **Names** field, you can specify a name for your print queue (where print jobs are sent), and the default directory to which you want the jobs sent (the **Spool Directory**

entry). You can probably leave these settings unchanged; note that giving a printer entry the name `lp` makes it the default printer.

Figure 5–5 Editing a Local Printer



The **Input Filter** contains information about your specific printer and its formatting requirements. To add your printer's information, click on the **Select** button. The **Configure Filter** dialog will open.

Once the **Configure Filter** dialog opens, your tasks are to perform the following:

- Pick an exact match for your make and model printer in the **Configure Filter** dialog. If you can't find an exact match, see if you can find a match that's close (such as the same manufacturer, but a different model). Select other options, such as paper size and resolution. Click on **OK** to accept your selections.
- Click on **OK** in the **Edit Local Printer Entry** dialog; you'll return to the **Print System Manager**.
- Highlight your new printer entry in the **Print System Manager** dialog, then restart `lpd` from the menu.
- Finally, with your printer highlighted in the **Print System Manager**, print a test page by selecting the option from the menu.

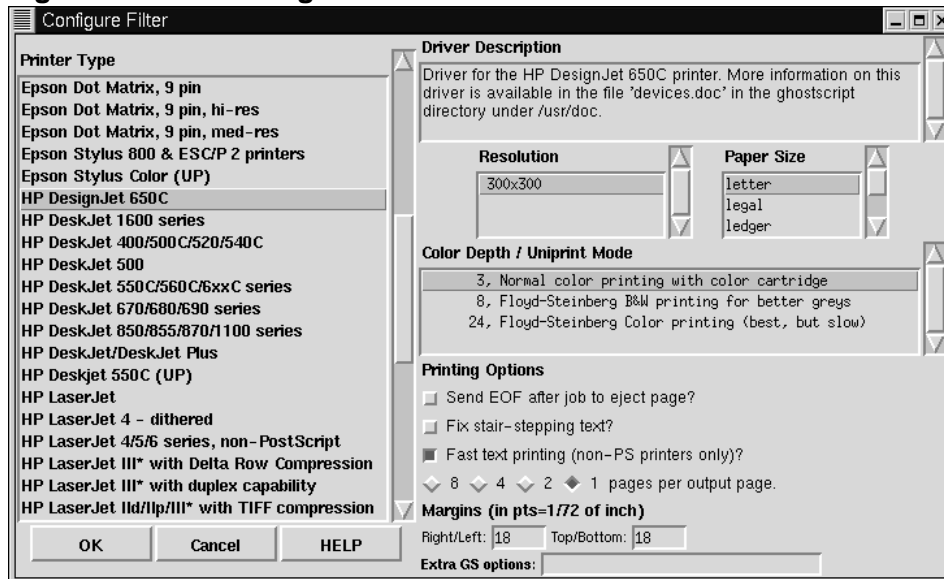
If you're unsuccessful at printing the test page -- or you need to modify the way it printed -- return to the **Configure Filter** dialog to select a different filter or change settings for your current filter.

If you're experiencing trouble, remember to check *The Hardware Compatibility List* and the *The Linux Printing HOWTO*, cited in the beginning of this chapter.

5.1.1 Selecting Input Filters

The **Configure Filter** dialog will look similar to Figure 5–6, *Selecting a Filter*. Select the appropriate filter for your printer. When you select a filter, you're selecting the correct way to convert file formats so that your printer can use them.

Figure 5–6 Selecting a Filter



The filter dialog is divided into several sections:

- **Printer Type:** Lists the available printer models for which you have drivers.
- **Driver Description:** Describes basic information about the driver, as well as the types of models for which the driver is appropriate.
- **Resolution and Paper Size:** Allow you to specify the quality of output to the printer, as well as the default paper size in the paper tray.
- **Color Depth/Uniprint Mode:** Allows you to specify additional resolution options for a color printer.
- **Printing Options:** Allows you to modify how jobs are printed.

In the **Printer Type** pane, at left, choose the exact make and model listed for your printer. If you can't find an exact match, try to find the closest match to your printer's make and model.

If you have questions about which filter may be appropriate for your printer, you'll find information about the driver in the box called **Driver Description**, at the right of the dialog.

Quite a few makes and models of printers are supported, but not all printers are fully supported -- or supported at all. In some cases, printers may require additional drivers; in other cases, the printer may be unsupported by Red Hat Linux. Again, refer to the Hardware Compatibility List at <http://www.redhat.com/support/hardware/> to verify that your printer is supported under Red Hat Linux.

If you find the manufacturer of your printer, but not your model number, you can check your printer's documentation to find out whether your printer can **emulate** any of the listed printers. You may also be able to find out whether a filter is suitable for your unlisted model by highlighting the filter in the left window and reading its description in the **Driver Description** window, at the right of the dialog.

Choose the resolution and default paper size for your printer. (Not all printer types will offer optional paper sizes.)

In the **Color Depth / Uniprint Mode** section, select the options you desire.

A little about Floyd-Steinberg

Floyd-Steinberg refers to a process of **dithering**, which is used to create shades or to smooth jagged lines.

5.1.2 Other Printing Options

The remainder of the **Configure Filter** dialog allows you to specify other printing options:

- **Send EOF after job to eject page** -- This option forces the printer to eject the page when the print job is finished. Uncheck this option if you receive an extra blank page after your print job.
-

- **Fix stair-stepping of text** -- Choose this option if your test page shows uneven lines of text, referred to as **stair-stepped**, which looks like:


This, for example,
is what stair-stepped
text might look like

- **Fast text printing (non-PS printers only)** -- If you don't have a PostScript printer, you can select this option, which will send text files directly to the printer, rather than send them to a program for conversion to PostScript, and then to Ghostscript to be printed. This option bypasses such a route for your text files, and speeds the process.
- **8, 4, 2, 1 pages per output page** -- This option refers to the command `mpage`, which allows you print multiple pages of plain text files on one page (either 2, 4, or 8). This can be useful for squeezing large printouts onto a single sheet of paper. Note, however, that this option is not available if you choose the "fast text printing" option.
- **Margins** -- This option specifically applies to the selection of multiple page printing. If you choose to print multiple pages (see the option above), you can specify the width of the margin on the page. If you choose to print single pages, this option will be unavailable to you.
- **Extra GS options** -- Refers to additional Ghostscript options which may be available for your printer. In general, you can leave this blank. Check in the **Driver Description** box for any options you might be required to add here.

5.1.3 Testing Your Printer Configuration

Once you've made your selections in the **Configure Filter** dialog, click on the **OK** button. You'll return to the **Edit Local Printer** dialog. Now, the **Input Filter** field will contain the information you'd chosen about your printer (see Figure 5-7, *Local Printer Entry with Filter*).

Figure 5–7 Local Printer Entry with Filter



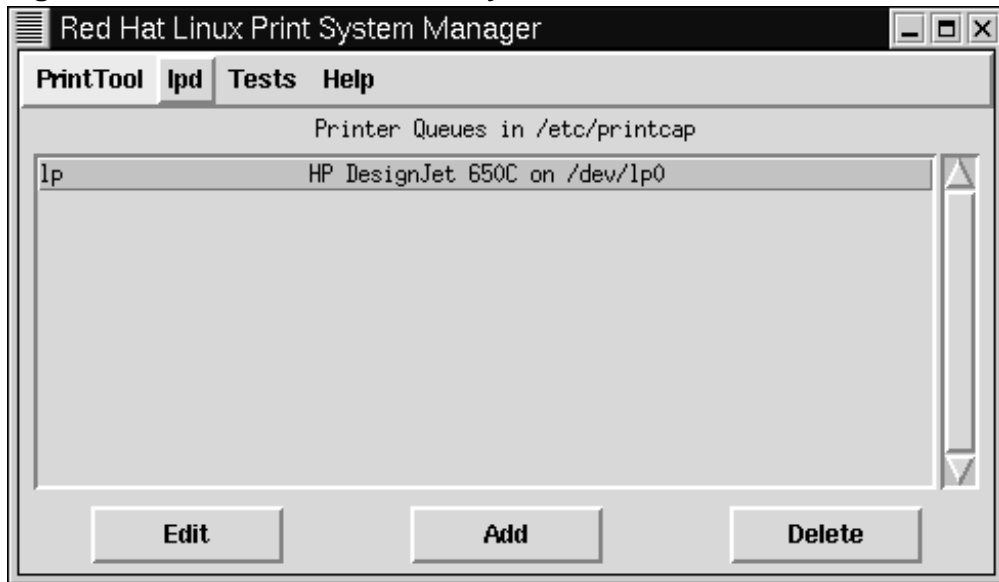
Leave the item **Suppress Headers** checked by default if you don't want extra data printed with your documents (such as the date and origination of the print job).

Click the **OK** button to continue; you'll return to the opening dialog of **Printtool**; this time, with your printer entry, as in Figure 5–8, *Your New Printer Entry*.

Now, it's time to test your setup.

- Restart `lpd` -- Restart the printer spooler daemon by selecting `lpd => Restart lpd` from the menu.
- Print a test page by highlighting your printer entry, then selecting `Tests => Print ASCII test page` (you can perform a test print on a PostScript test page, as well).

Figure 5–8 Your New Printer Entry



If you want to make modifications to your printer selection, highlight the printer and select **Edit** at the bottom of the **Printtool** window.

When you're satisfied with your configuration, you can exit by selecting **Quit** from the **PrintTool** menu, or by clicking on the **x** in the upper right of the application's window.

Once you return to your prompt, you can exit from the root login by clicking on the **x** in the upper right of the Xterm window, or by typing `exit` at the prompt. Then, type `exit` again to close the Xterm.

5.2 Printing Utilities in KDE and GNOME

Once you've configured your printer, you'll find KDE and GNOME provide utilities to help you monitor and modify printer settings and printing jobs.

5.2.1 Using Ljet Tool in KDE

In KDE, Ljet Tool can help you configure a Hewlett-Packard LaserJet printer in KDE.

Also called **KLJetTool**, the utility can allow you to modify your printer settings for all KDE-compliant applications. **Ljet Tool** can be used by regular users; not just by the root user.

Possible conflicts

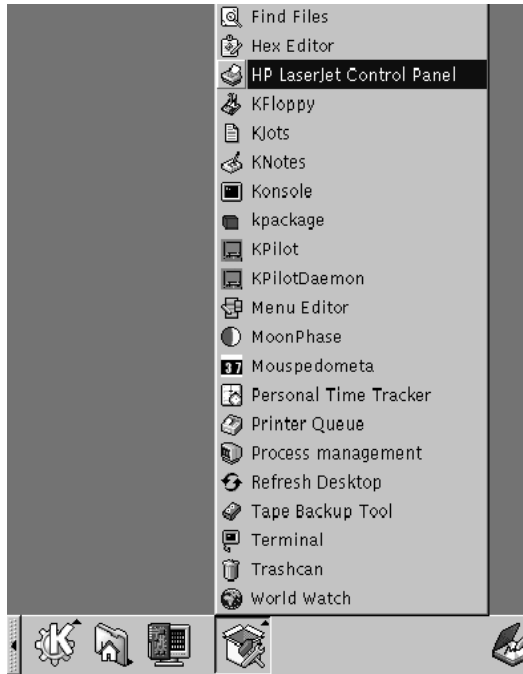
Ljet Tool is best suited for Hewlett-Packard LaserJet and compatible printers. The utility uses PCL, the same printer control language to control options used by Hewlett-Packard. If your printer isn't a Hewlett-Packard LaserJet or PCL-compatible printer, you may experience problems if you try to use **Ljet**.

5.2.2 Starting Ljet Tool

You can start **Ljet** from the **Panel** or from the shell prompt.


To start the utility from the shell prompt, open an Xterm window and, at the prompt, type `kljettool`.

To start **Ljet** from the **Panel**, go to the **Utilities** menu, then go to **HP LaserJet Control Panel** (see Figure 5-9, *Starting Ljet from the Panel* for an example).

Figure 5–9 Starting Ljet from the Panel

In four tabbed dialogs, the utility allows you to specify details like paper size, resolution, font size and more.

Most of the options are fairly self-explanatory, but some of the settings in the **Fonts** dialog may be a little confusing: specifically, the settings for **Language** and **Symbol Set**. See Figure 5–10, *The Fonts dialog in Ljet Tool* for an example.

Figure 5–10 The Fonts dialog in Ljet Tool

The above dialog allows you to choose default settings for the characters which are used to print documents.

The **Language** entry defines which printer language is used to process your documents. By default, this is set to **PCL**, which is the standard printer language for Hewlett-Packard printers. You can use the drop-down list to select PostScript, however, if your printer has such capabilities.

The **Symbol Set** entry refers to the default character set you'd prefer to use when printing. The setting of **PC8** is the default mode; if you want to print special math, Windows or non-English characters, however, you can make your selections here.

5.2.3 The Printer Queue in KDE

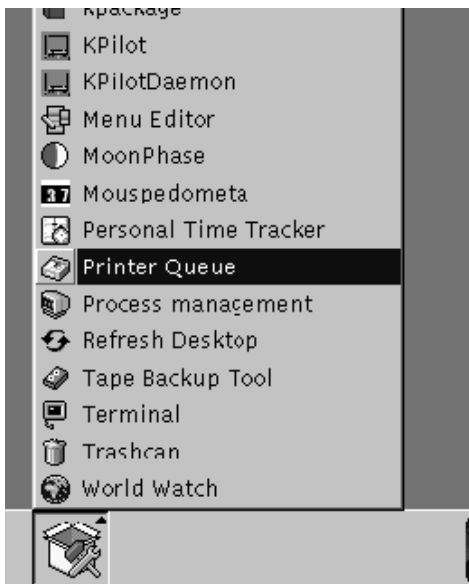
On the KDE desktop, the printer icon represents your printer queue. You can drag and drop files onto the printer icon to send them to the printer.

You can open the printer queue either by clicking on the printer icon on your desktop or by going to the **Printer Queue** item on the **Utilities** menu of the **Panel** (see Figure 5–11, *Opening the Printer Queue from the Panel*).

You can also open the printer queue by typing `klpq` at the shell prompt.

When you first start the printer queue application, you'll be asked which print spooler you use. By default, **BSD** will be checked, and you probably won't have to alter this setting.

Figure 5–11 Opening the Printer Queue from the Panel



When you open the printer queue, you won't see any entries in the queue (see Figure 5–12, *The Printer Queue Window*), unless you're currently printing a document.


Figure 5–12 The Printer Queue Window



A document you send to the printer will be displayed in the Printer Queue utility. See Figure 5–13, *A Print Job in the Queue* for an example of a document listed in the queue.

Once in the queue, you can see details about the job, such as the user who sent the job to the printer, its order in the printer queue -- if there is more than one job -- the size of the document and more.

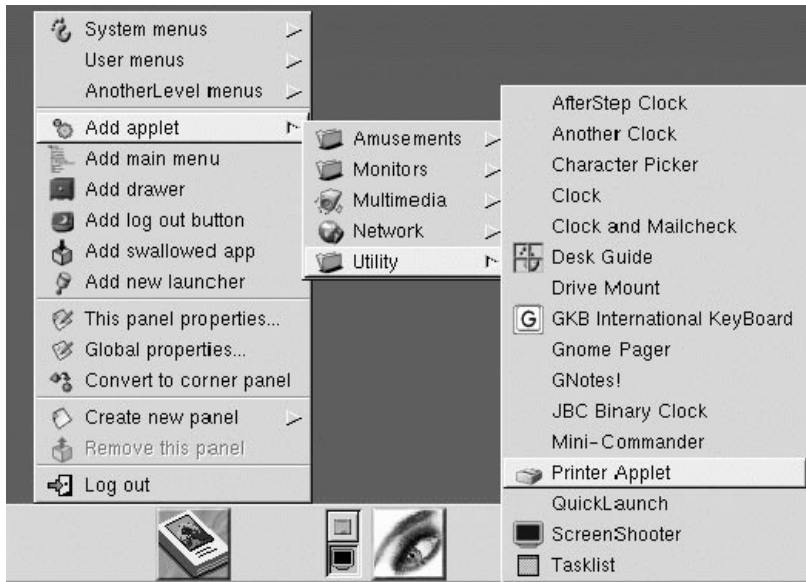
The Printer Queue also allows you to change the order of jobs, or to remove printing jobs.

Figure 5–13 A Print Job in the Queue

5.2.4 The Printer Applet in GNOME

GNOME offers a convenient way to drag and drop documents to the printer with the Printer Applet. The utility can be placed on the **Panel**.

Figure 5–14 Adding the Printer Applet to the Panel



To add the applet to the **Panel**, right-click on the **Panel** and drag your cursor to **Add Applet** => **Utility** => **Printer Applet** (see Figure 5–14, *Adding the Printer Applet to the Panel*). Release your mouse button over **Printer Applet**; the applet will appear on your




Panel and look like:

To use the applet, drag a document from a folder or from the file manager and "drop" it over the **Print** icon on the **Panel**. (When you hover your document over the icon and an outline surrounds the icon, you'll know that you can drop your document on the applet.) After a few moments, your document should begin printing.

You can adjust a couple of the settings in the **Printer Applet** by right-clicking on the applet in the **Panel** and selecting the **Properties** dialog.

In **Printer properties** (see Figure 5–15, *Properties in the Printer Applet*), you can select a name for the printer label on the **Panel**. Your default print command will most likely be `lpr`, so you should leave the **Print command** entry the default `lpr`.

Figure 5–15 Properties in the Printer Applet



6 Getting Online

Normally, **PPP** accounts are used to dial in to an Internet Service Provider (**ISP**). Using PPP, which stands for **Point-to-Point Protocol**, you connect to a large network (the Internet) through your ISP; your machine becomes a part of that network and you use the resources of the network.

Both GNOME and KDE offer useful tools for establishing and modifying PPP accounts to ISPs. These utilities remove much of the guesswork in creating an online connection; in many cases, your only requirement is a valid Internet account.

The utilities -- **RP3** for GNOME and **Kppp** for KDE -- can be interchangeable, too. You can use **RP3** to set up and connect in KDE; similarly, you can use **Kppp** if you're using GNOME.

If you use **RP3** or **Kppp** outside of their respective environments, however, you will lose the ability to **dock** the running applications. That means you won't be able to integrate **RP3** into the KDE **Panel**, for example -- or **Kppp** in the GNOME **Panel**. When they're docked on their **Panels**, both utilities provide information about your connection, such as connection time, speed of the connection and more. You will be able to **minimize** these utilities on the **Panel**, however.

Note, also, that if you create an account in GNOME with **RP3**, but would rather use **Kppp**, you'll have to set up your account again in that application; in other words, information you enter in one application won't be available to the other application.

You must be root

Because you'll be making changes to your system, you must be root in order to create a PPP account. If you're in your user account, you'll be asked for your **root password** before proceeding in either **RP3** or **Kppp**.

Your own ISP may have specific connection requirements for their service which differ from the instructions in this chapter. Before connecting with either of these tools,

check first with your ISP for any specific instructions that they provide, including the following information:

- The phone number that your machine must dial into at your ISP.
- Your login name and password for the Internet account.
- Possible **gateway address**: Some ISPs may require you to specifically type in the address for the ISP's gateway.
- **DNS entries**: DNS means **Domain Name Service**. The DNS tracks **IP** (Internet Protocol) addresses; each computer connected to the Internet must have an IP address, which is a unique set of numbers like *2xx.2xx.2x.2*. You may receive one or more DNS entries from your ISP when you sign up. Both RP3 and Kppp provide spaces for you to enter these numbers.

6.1 Using RP3 to Connect

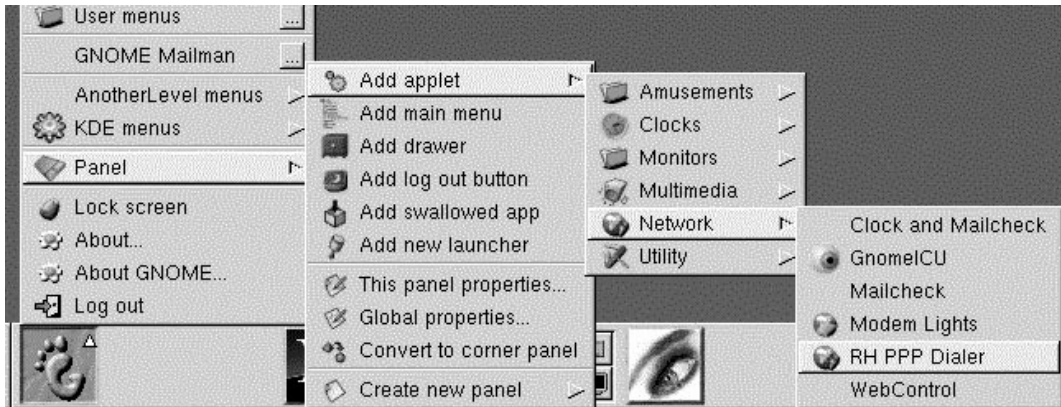
RP3, the Red Hat PPP Dialer utility, can make easy work of connecting to the Internet or to some other network devices. You can use RP3 to connect your modem to a PPP account, and then you can view statistics about the connection, such as the time, speed and cost of the current connection.

Don't mix utilities

While RP3, the **Dialup Configuration Tool**, is useful for creating a PPP account, you should avoid using the application in conjunction with another network configuration utility, called **Network Configurator** (`netcfg`). Both applications can create a PPP dialup connection, but you should use one to the exclusion of the other. For most account purposes, the **Dialup Configuration Tool** will work well.

You can start RP3 by going to the GNOME **Main Menu Button** => **Panel** => **Add applet** => **Network** => **RH PPP Dialer** (see Figure 6–1, *Adding RP3 to the Panel*). The application will then dock into the **Panel** and remain there until you remove it (by right-clicking on the docked application and selecting **Remove from panel**).

Figure 6–1 Adding RP3 to 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 modems and Ethernet cards, 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 onto the **Panel**.

RP3's **Dialup Configuration Tool** dialog (see Figure 6–2, *Finding the Dialup Configuration Tool on the Menu*) allows you to easily configure your system for an Internet connection, helping you to locate your modem, enter the account information and connect.

Figure 6–2 Finding the Dialup Configuration Tool on the Menu

6.1.1 Creating a New Internet Account

To create a new PPP account, start RP3. Go to **Main Menu Button** => **Internet** => **Dialup Configuration Tool**.

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.

If you don't already have devices configured, you'll see the **Add New Internet Connection** dialog (see Figure 6–3, *The Add New Internet Connection Dialog*).

Figure 6–3 The Add New Internet Connection Dialog

If devices are detected

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. Once the application is running, right-click on the application to open the options menu and select **Configure PPP**.

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 during 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 (as in Figure 6–4, *Searching for a Modem*).

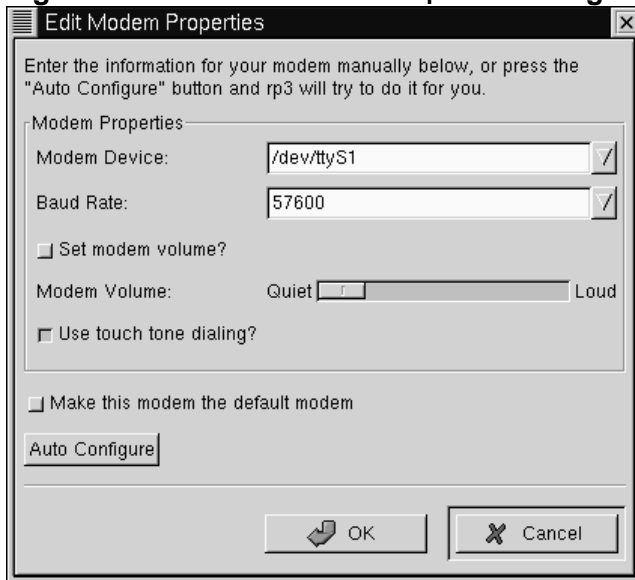
Figure 6–4 Searching for a Modem

If your modem isn't detected during the search, you'll receive a message stating that it couldn't be found.

You can still enter details about your modem's location, speed and other information, in the **Edit Modem Properties** dialog, as in Figure 6–5, *The Edit Modem Properties Dialog*. This dialog will appear if no modems were detected on your system. Refer to your modem manufacturer's documentation for specifics to enter here.

What are Winmodems?

If your machine uses a modem which relies on software drivers for some of its functionality, you've probably got a "Winmodem." Because they depend on Windows-based software to operate correctly, these types of modems don't function in Linux. To find out more about this subject, go to <http://www.o2.net/~gromitkc/winmodem.html>

Figure 6–5 The Edit Modem Properties Dialog

In the **Edit Modem Properties** dialog, you can specify the location of your modem, its baud rate, your preference regarding the speaker volume during a call and whether to use Touch Tone dialing.

Although **RP3** attempts to detect the correct port, you can adjust the settings here.

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 port to which your modem is connected. Then, in **Baud Rate**, enter the fastest recommended speed for your

modem. Due to most modems' data compression capabilities, select a baud rate that is faster than the modem's advertised speed (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.

Check **Make this modem the default modem** to use your current modem settings each time you run RP3.

You can also rescan your system for modem information by selecting the **Auto Configure** button, then choose to keep the settings it finds.

Once you're satisfied with your settings, click on **Next** to proceed to the following screen.

Tip

Even if your modem is automatically detected, you can adjust the modem's settings by choosing to **Manually create a modem** after RP3 detects your device. Your choices, such as location of modem, baud rate and volume, will be identical to those you'd find if the modem had not been detected.

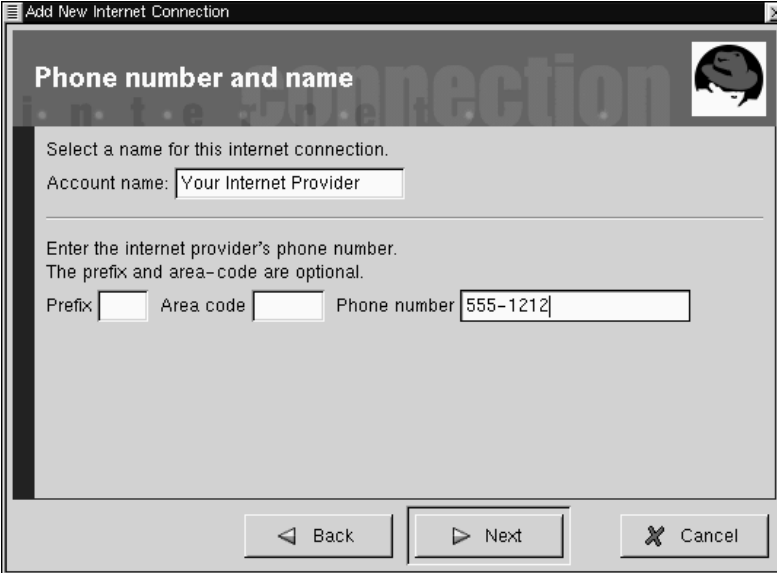
In the **Phone number and name** dialog (see Figure 6-6, *Adding a Name and Phone Number*), you should enter both the name you would like to give your account and the telephone number of your ISP. Usually, it isn't necessary to provide the prefix or area code for your ISP.

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.

To learn about other options

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 this document a "page" at a time; use the [B] key to move back, and use the [Q] key to quit.)

Figure 6–6 Adding a Name and Phone Number



The screenshot shows a window titled "Add New Internet Connection" with a globe icon in the top right corner. The main heading is "Phone number and name". Below this, there are three sections of text and input fields:

- Text: "Select a name for this internet connection."
- Text: "Account name:" followed by a text box containing "Your Internet Provider".
- Text: "Enter the internet provider's phone number. The prefix and area-code are optional."
- Text: "Prefix" followed by an empty text box, "Area code" followed by an empty text box, and "Phone number" followed by a text box containing "555-1212".

At the bottom of the dialog, there are three buttons: "Back" (with a left arrow), "Next" (with a right arrow), and "Cancel" (with an 'X' icon).


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 (see Figure 6–7, *Specifying a User Name and Password*).

Figure 6–7 Specifying a User Name and Password

The image shows a Windows-style dialog box titled "Add New Internet Connection". The main heading is "User name and password". Below the heading, there are two input fields. The first is labeled "User name:" and contains the text "yourusername". The second is labeled "Password:" and contains a series of asterisks "*****". At the bottom of the dialog, there are three buttons: "Back" with a left-pointing arrow, "Next" with a right-pointing arrow, and "Cancel" with an 'X' icon.

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 6–8, *Checking Your Entries*).

Figure 6–8 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/wvdial.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 6.1.4, *Customizing and Configuring* for more information on customizing the settings.

6.1.2 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.

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 from your session by left-clicking on the monitor, then choosing **Yes** to close the interface. (Also, you can use the **Disconnect from...** entry by right-clicking on the monitor.)

6.1.3 Working with the RP3 Monitor

When the RP3 monitor is docked on the **Panel**, it will look like Figure 6–9, *The RP3 Monitor Docked (Second Icon From Left)*.

Figure 6–9 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 6–10, *The RP3 Monitor*).

Figure 6–10 The RP3 Monitor



Beneath the graphs, you'll find text entries which will display the bytes received per second and either your total connection time or the total cost of your current session (if you've specified this option when you configured the account).

6.1.4 Customizing and Configuring

RP3 makes easy work of setting up an Internet account, but you may still find you've got a little tweaking to do before you're up and running.


By default, when RP3 connects to your ISP, it starts `pppd`. What is `pppd`? It means **Point-to-Point Protocol daemon**; briefly, it helps your machine to "shake hands" with your ISP.

To learn more about `pppd`

You can learn more about `pppd` by reading its man page. At the shell prompt, type `man pppd`.

You can customize and configure account information by right-clicking on the RP3 monitor and selecting the **Configure PPP** entry.

Figure 6–11 The Internet Connections Dialog



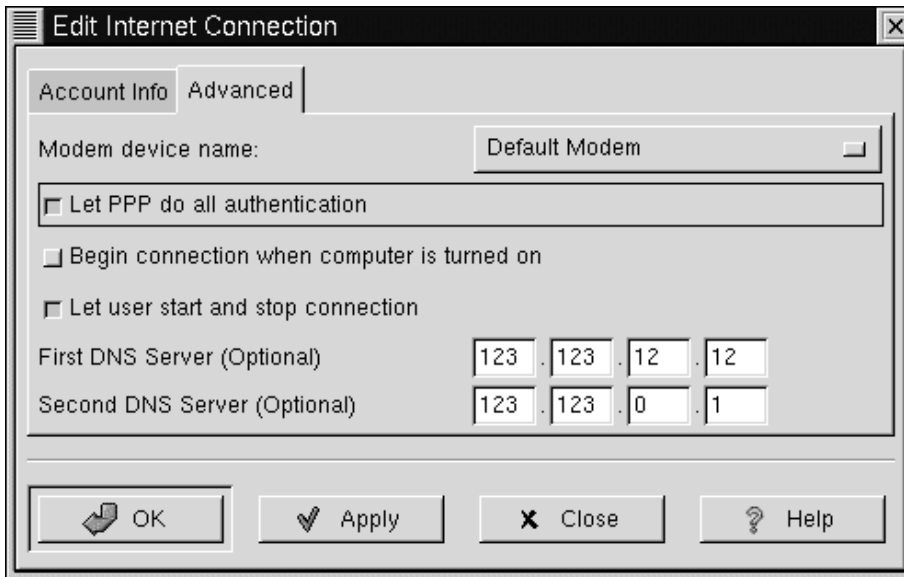
In the **Internet Connections** dialog which appears, you can add a new account, delete an account and modify current account configuration. You can also **debug** the account, meaning that you can test your settings to see if there are any problems.

Selecting the **Add** button will begin the **Configure PPP** tool, allowing you to enter new account information.

If you choose to modify details in a current account, select the **Edit** button. When the **Edit Internet Connection** dialog opens, you can change account details in the **Account Info** tab. Among the fields you can modify are your login name, your password, and your ISP's telephone number.

If your ISP provided you with **DNS numbers**, which may be necessary for a successful connection, you can enter them in the **Advanced** tab of this dialog (see Figure 6–12, *The Advanced Settings Dialog* for an example).

Figure 6–12 The Advanced Settings Dialog



Among other options in the **Advanced** tab, you can select whether you want people who have user accounts on your system to be able to connect and disconnect. You can also choose whether you want to start your connection when your machine starts (**Begin connection when computer is turned on**).

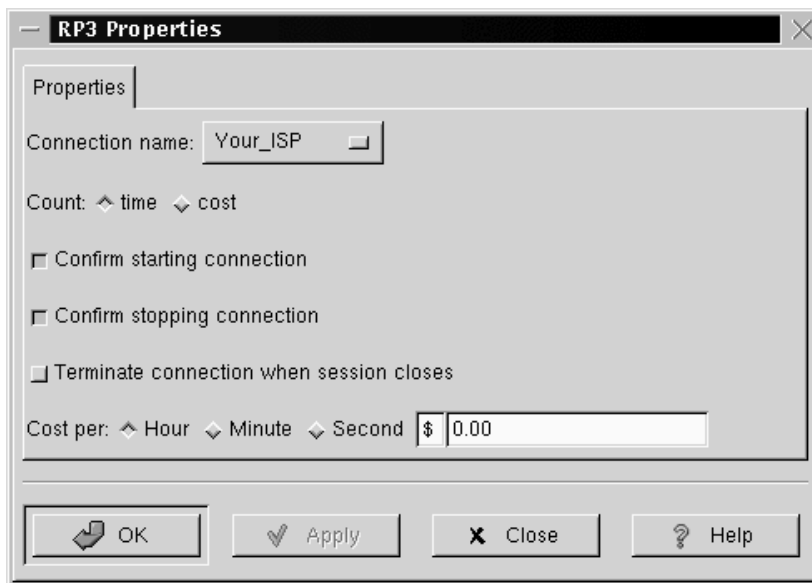
The option **Let PPP do all authentication** will cause `pppd` to start immediately upon a successful connection. Learn more about this a little later, in Section 6.1.5, *Troubleshooting*.

In the **Modems** tab of the **Internet Connections** 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.

You can determine how you want to view some details about your connection status. To change the default settings, right-click on the application while it's running and, in the menu which appears, left-click once on **Properties**.

In the **Properties** dialog (see Figure 6–13, *The Properties Dialog*), you'll find the following entries:

Figure 6–13 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:** If you pay for Internet access by the hour, minute or second, you can specify the cost for your access here. When you're connected, your connection costs will then be reflected on the **RP3** monitor.


6.1.5 Troubleshooting

Sometimes, things may not go as well as you'd hoped when you're trying to establish a connection. That's when **debugging** can come in handy. Briefly, debugging allows you to see a little of what's going on behind the curtain when you're trying to connect.

To start, launch **RP3** and open the **Accounts** tab in the **Internet Connections** dialog. Now, highlight the account you want to debug, and select the **Debug** button in the options to the right.

Once you choose to debug the account, you'll see a window similar to Figure 6–14, *Debugging an Account*, showing you the commands which are used to initialize your modem, connect to your ISP and perform a login to your Internet account.

Figure 6–14 Debugging an Account



The lines may slip by too quickly for you to catch anything unusual. Not to worry: You can scroll through the window to study the information and, for easier review, you can choose the **Save Log** button to save the contents of the window as a text file, which you can print or open in an editor or word processor.

You may find, for example, that you can connect to your ISP, but you're not correctly logging in. Check first to make certain that you entered the correct user name and password for your account.

It's also possible that the `pppd` is starting too early, before you have a chance to log in with your name and password. In this case, you *appear* to be successfully connected, yet when you try to view a Web page with Netscape Navigator, you're met with error messages.

If this is the case, stop debugging the connection by selecting **Close** or **Cancel**, then return to the **Edit** feature in the **Accounts** tab.

In the **Edit Internet Connection** dialog, select the **Advanced** tab and deselect the option **Let PPP do all authentication**. (By default, this option isn't selected when you create an account.) If everything else looks correct, choose **OK** to return to the **Accounts** tab.

Now, debug the connection again. You might find that your login procedure goes a little more smoothly, since `pppd` isn't started immediately, but rather waits until your login name and password are sent to your ISP.

Working with the debug feature can help you whittle down whatever problem might be hindering you from a working connection. When you're finally successful, close the window to stop debugging. Since the debug option is only used to diagnose your setup, it won't maintain your connection to your ISP.

When you've successfully configured your connection, start the interface as described in the beginning of this section.

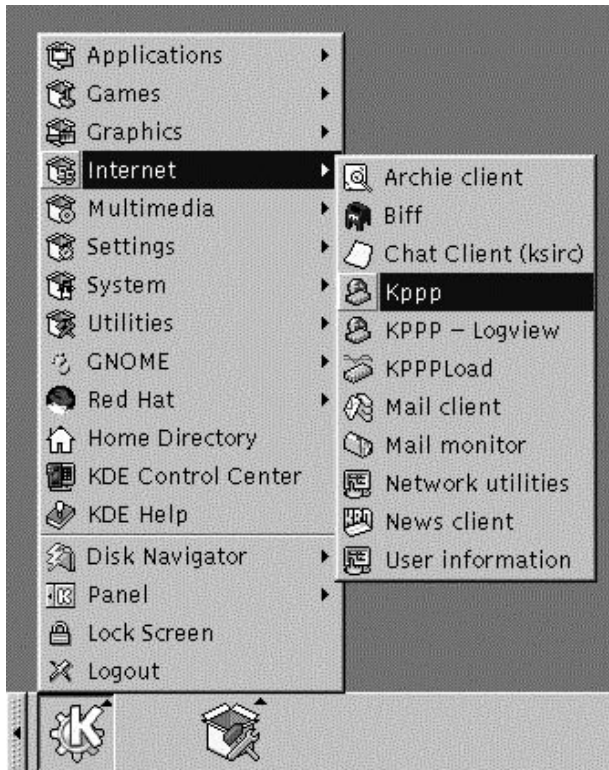
6.2 Connecting with Kppp

Kppp is one of the useful utilities included in KDE. Kppp can make establishing your Internet connection easy -- and easy to modify. A dialer and configuration application, Kppp works with `pppd` -- the Point-to-Point Protocol daemon.


6.2.1 Starting Kppp

To start Kppp, go to the **Main Menu K** icon => **Internet** => **Kppp** (see Figure 6-15, *Starting Kppp from the Menu*).

Figure 6–15 Starting Kppp from the Menu




Type the password for the root account when it's requested. The opening dialog for Kppp will be "grayed out," because you haven't created an account yet (see Figure 6–16, *The Main Dialog of Kppp*).

Figure 6–16 The Main Dialog of Kppp

6.2.2 Configuring Kppp

To begin, select the **Setup** button. The next dialog to open will be the Kppp configuration dialog (similar to Figure 6–17, *The Kppp Configuration Dialog*).

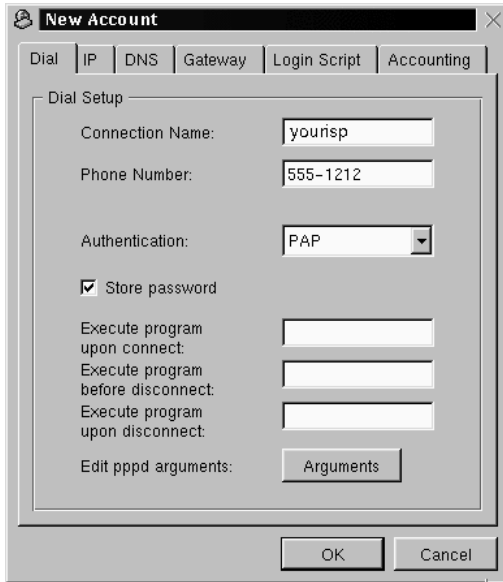
Figure 6–17 The Kppp Configuration Dialog



Click the **New** button to enter your Internet account information.

The Dial Tab

In the **New Account** dialog, you can specify the account details, such as name, phone number, DNS entries and other information. The information is divided into "tabbed" dialogs (see Figure 6–18, *The New Account Dial Tab*).

Figure 6–18 The New Account Dial Tab

In the **Dial** tab, enter the following:

- **Connection Name:** The name you wish to give your account (for example, Mind-spring or My ISP).
- **Phone Number:** The phone number to connect to your ISP.
- **Authentication:** From this drop-down list, you can select the following methods of login authentication: PAP, Script-based, Terminal-based, or CHAP. (These are methods used to verify your login. You may want to check with your ISP to determine which type of authentication method you must use.) If you're required to use a script-based login, *The Login Script Tab* in Section 6.2.2 can assist you in entering your script.
- **Store password:** Check this option if you don't want to be required to enter your password each time you want to log on.

The remaining fields allow you to specify a program you want to run when your connection is established (or after you've disconnected).


You can also specify options you want to pass to `pppd` in the **Edit pppd arguments** by pressing the **Arguments** button.

Unless you know what you're doing, however, you should leave these additional options alone.

Static or Dynamic IP

In the **IP** tab, you should specify whether you'll be using **dynamic IP addresses** or, if your ISP has provided you with your own IP address you can enter it here (see Figure 6–19, *Specify IP Addressing*).

Figure 6–19 Specify IP Addressing

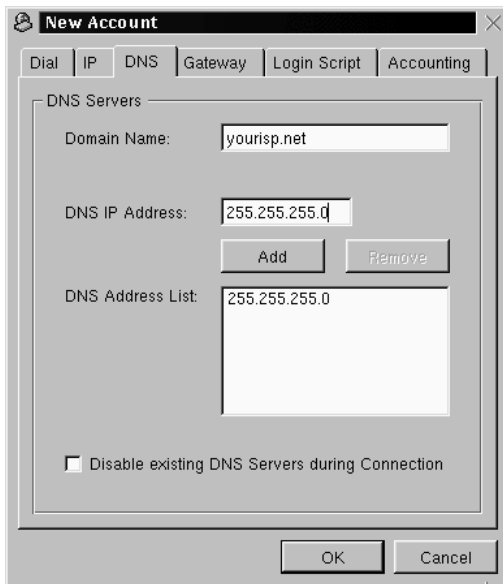


(Hint: If your ISP has *not* provided you with an IP address, you should check **Dynamic IP Address**.)

DNS entries


In the **DNS** tab (see Figure 6–20, *Specify Domain Name and DNS Entries*), you should enter the name of your ISP's domain name (such as *earthlink.net*).

Figure 6–20 Specify Domain Name and DNS Entries



If your ISP provided you with DNS addresses, enter them in the **DNS IP Address** field. When you've finished typing in the numbers, click on the **Add** button, and the numbers will be entered into the **DNS Address List** box. (By default, you won't be able to click on the **Add** button until you've entered the entire DNS number in the address field.)

Note that the field **Disable existing DNS Servers during Connection** will override the DNS list. If you've entered DNS numbers, you should leave this option unchecked.

Figure 6–21 Select Default or Static Gateway

Specify Gateways

In the next tab, **Gateway** (see Figure 6–21, *Select Default or Static Gateway*), select either **Default Gateway** or **Static Gateway**. If your ISP has specified an IP address you should use as a static gateway, enter that IP number in the space provided.

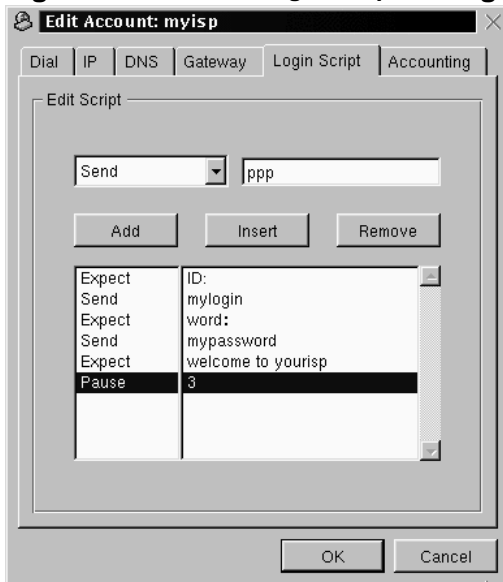
Assigning the "default route"

Unless you know what you're doing, you should leave **Assign the Default Route to this Gateway** box checked. Unchecking this option could cause problems for your software when you're online.

The Login Script Tab

The **Login Script** tab allows you to specify how you want to login to your ISP (see Figure 6–22, *The Login Script Dialog*). If you chose to use a script-based login in the **Authentication** procedure of the **Dial** tab, you can enter that script here.

Figure 6–22 The Login Script Dialog



Not all ISPs require a script-based login procedure; you should check with your ISP to find out specific entries which might be necessary in a login script.

A login script helps your machine negotiate a successful connection with your ISP. Your machine *expects* a string of information from your ISP, then *sends* a string of information or prompts. In this way, your machine "shakes hands" with your ISP.

In the **Login Script** tab, use the drop-down list at the left of the dialog to create a series of steps you want your machine to follow when you log in to your ISP. You can select a variety of options to create your script, step by step. In the window to the right of the drop-down list, type in the specific information you want to supply with the option from the drop-down list.

Click the **Add** button to add your step. If you must modify your script, you can choose an option, supply the information and use the **Insert** button to place it in the script, rather than having to recreate the script from scratch.

You can also remove a step in your script by highlighting the entry and clicking the **Remove** button.

If you have difficulty connecting to your ISP after you configure Kppp, check the option **Show Log Window** on the main Kppp dialog. The log window will open when you try to connect to your ISP, showing you what your ISP is expecting and what your machine is sending.

In Figure 6–22, *The Login Script Dialog*, for example, the login script is for an ISP which sends a prompt for the login ID, then for the password and finally, after a brief pause, requires that a user send the string **ppp**. The script looks like:

Table 6–1 A Sample Login Script

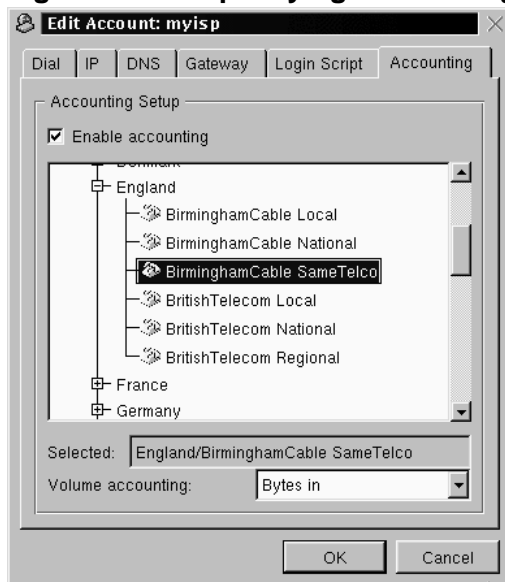
Option	Enter	Purpose
Expect	ID:	Wait for the ID: login prompt from the server
Send	mylogin	Send your login name for the account
Expect	word:	Wait for the password prompt from the server
Send	mypass- word	Send the account's password
Expect	welcome to yourisp	Wait for the welcome message from the server
Pause	3	Specifies a pause of three seconds (a value for Pause will be in seconds, by default)

It may not be necessary for you to use a script-based login procedure, however; this depends on requirements from your ISP. In some cases, using PAP or CHAP may be all that's necessary. Again, check with your ISP for specific login requirements.

The Accounting Tab

In the **Accounting** tab, you can specify whether and how you want Kppp to monitor your online activities (see Figure 6–23, *Specifying Accounting*).

Figure 6–23 Specifying Accounting



Accounting can be useful if your ISP bills you by the number of bytes transferred during your time online -- rather than by a flat monthly rate, for example. Referred to as **volume accounting**, this tool will count the number of bytes to or from your machine on the Internet.

Once you're finished creating your account, click on the **OK** button.

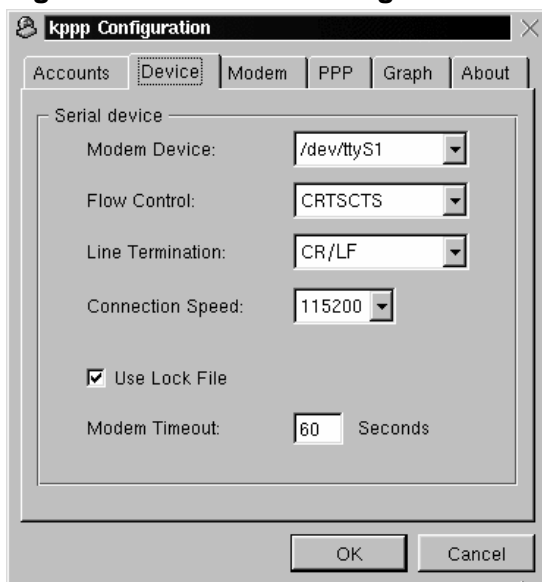
Next, in the **kppp Configuration** window, you'll see your newly created account. (If you want to modify the account, highlight the account by clicking on it once, then selecting the **Edit** button to the side of the **Account Setup** window.)

Specifying a Device

Now, it's time to provide the location and speed of your modem, in the **Device** tab (see Figure 6–24, *Device Configuration*). Here, you specify where your modem is located, its baud rate, some hardware settings and other specifics.

In **Modem Device**, you'll find a drop-down box allowing you to choose the location of your modem.

Figure 6–24 Device Configuration



The following list is a comparison of serial line names in MS-DOS (and Windows) and Linux. If you know that your modem is attached to COM2 under Windows, for example, its location in Linux will be `/dev/ttyS1`.

- COM1 = `/dev/ttyS0`
- COM2 = `/dev/ttyS1`
- COM3 = `/dev/ttyS2`
- COM4 = `/dev/ttyS3`

Once you've selected the location of your modem, you can select the **flow control**.

Although you have a choice in the drop-down list for **Flow Control**, the default -- **CRTSCTS**, which specifies **hardware flow control** -- is the recommended setting (**XON/XOFF** specifies **software flow control**).

The following entry, **Line Termination**, allows you to specify a correct setting for your modem. Often, the **CR/LF** entry is correct for most modems. If, however, you have difficulty **querying your modem** (see *Testing Your Modem* in Section 6.2.2), this may be one area to which you should return to adjust settings.

The **Connection Speed** entry allows you to select the highest baud rate for your modem (for example, a 56k modem can typically be set to 115200).

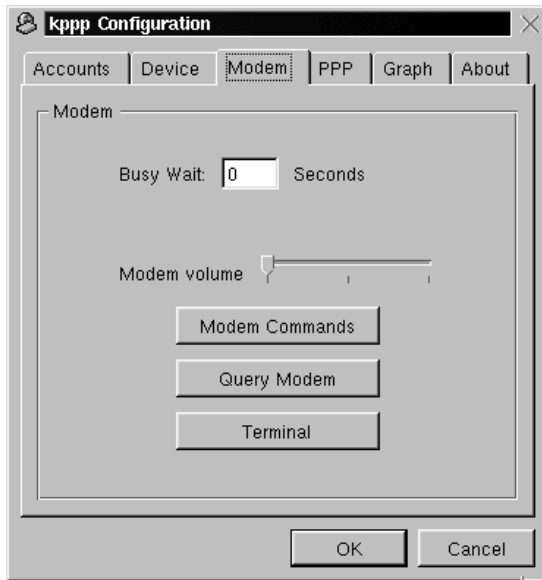
By default, **Use Lock File** is checked. This means that Kppp will **lock the device** while you're online -- preventing others from using the modem.

The **Modem Timeout** entry provides the time, in seconds during which Kppp waits for a CONNECT response after dialing.

Testing Your Modem

In the **Modem** section, you can test your modem's responsiveness to your configuration (see Figure 6–25, *The Modem Tab*). By adjusting the slider next to **Modem volume**, you can modify the volume of your modem while you're dialing and connecting.

Figure 6–25 The Modem Tab

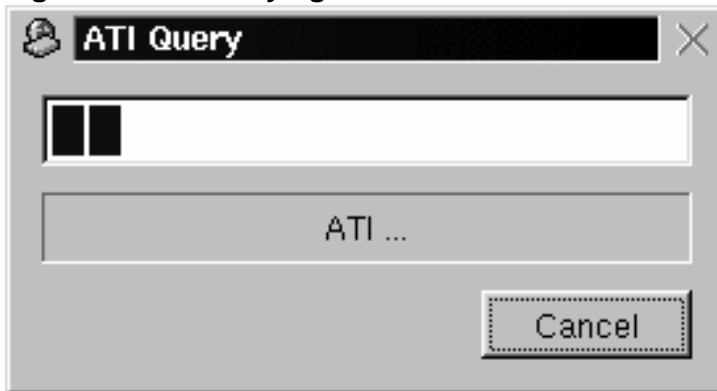


When you select the **Modem Commands** button, you'll be presented with the **Edit Modem Commands** dialog, which will allow you to modify how your modem dials, connects and more.

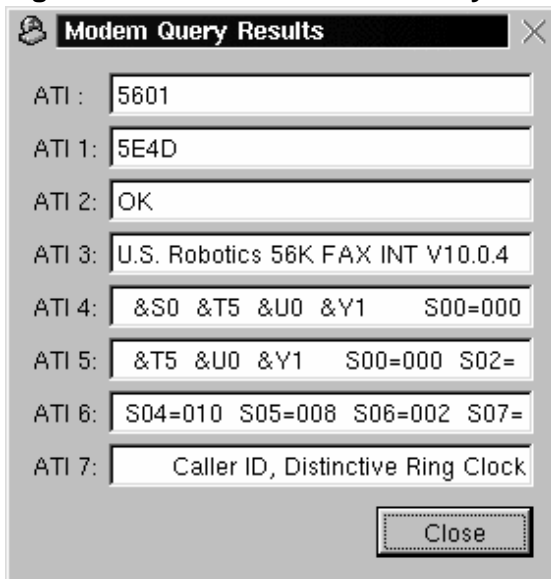
Check your modem's documentation

Check the documentation supplied with your modem for specifics on your modem's **initialization string**, **dial string** and other settings.

Click on the **Query Modem** button to have Kppp search your system to identify the modem. A window similar to Figure 6–26, *Querying for the Modem* will open to show the progress of the query.

Figure 6–26 Querying for the Modem

If the query was successful, you'll see the results in the dialog which follows, showing details which Kppp has gleaned from your modem (see Figure 6–27, *Results of the Query*).

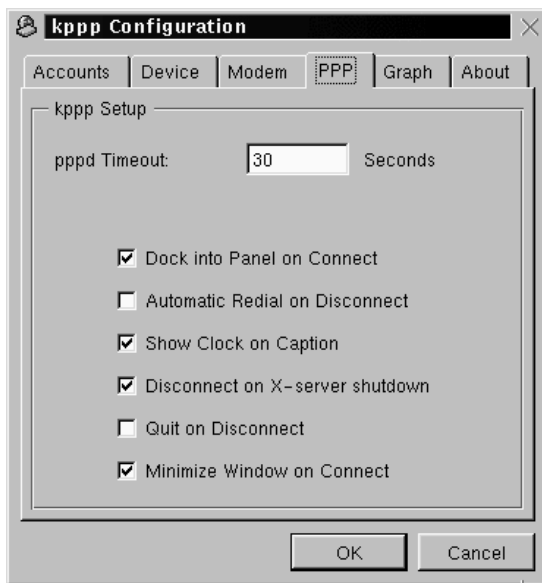
Figure 6–27 Results of the Query

If you select the **Terminal** button, Kppp will open a **mini-terminal**, with which you can test your modem configuration.

The PPP and Graph Tabs

In the **PPP** tab, you can determine such options as automatic reconnection or disconnection, display of Kppp on connection and more.

Figure 6–28 The PPP Tab



You can select from the following options:

- **pppd Timeout:** Specifies the time allowed to establish a connection, before giving up.
- **Dock into Panel on Connect:** Select this to have Kppp displayed as an icon in the Panel once you establish a connection. (This option overrides the **Minimize Window on Connect** selection.)
- **Automatic Redial on Disconnect:** Place a checkmark in this box if you want to have Kppp immediately redial your service if you're disconnected.

- **Show Clock on Caption:** Check this to see a clock displaying your connection time.
- **Disconnect on X-server shutdown:** If you choose this option, your connection will be terminated when your X server shuts down. This is a useful option if you're likely to shutdown your X session without taking the time to disconnect, because it causes a disconnection when you log out.
- **Quit on Disconnect:** When you've disconnected from your ISP, Kppp will also quit. Otherwise, you'll return to the initial Kppp window.
- **Minimize Window on Connect:** This will minimize Kppp to the taskbar.

In the **Graph** tab, you can modify the colors of the **throughput graph**, which is a visual chart showing you the bytes exchanged between your machine and your ISP while you're online.

The **About** tab displays some information about Kppp, such as its authors, copyright and credits.

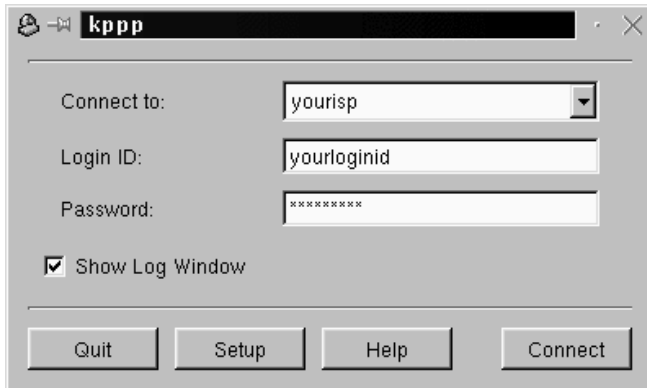
6.2.3 Connecting and Disconnecting with Kppp

Once you've configured Kppp, you're ready to try it out.

From the **Main Menu K** icon, go to **Internet** and click on **Kppp**. Before the application starts, you'll be required to enter your root password. (Note that you can remain in your user account, however; it isn't necessary to log in as root.)

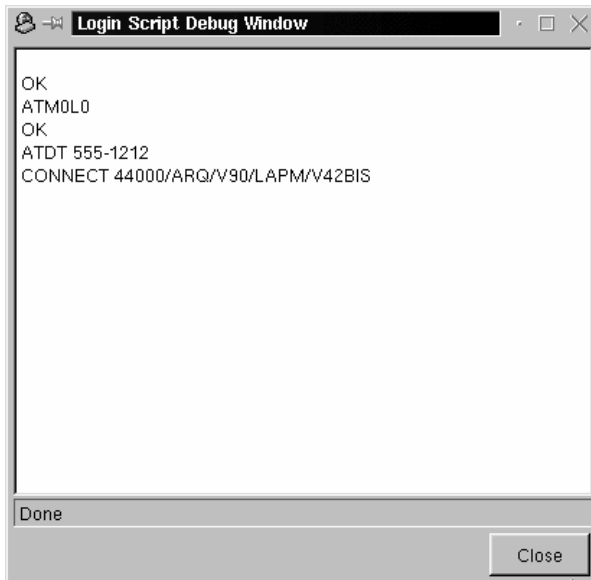
In the main dialog for Kppp which opens, enter the login ID and password for your Internet account (as in Figure 6–29, *The Main Dialog for Kppp*).

Figure 6–29 The Main Dialog for Kppp




If you want to **debug** the configuration, or watch the progress of your connection, check the box marked **Show Log Window**.

Figure 6–30 The Login Script Debug Window in Kppp



To begin your session, click on the **Connect** button. A window, similar to Figure 6–31, *The Progress Window*, will show the status of your connection, such as the initialization of your modem, and the progress of the transaction. (You can also open the log window by clicking on the **Log** button in this progress window.)

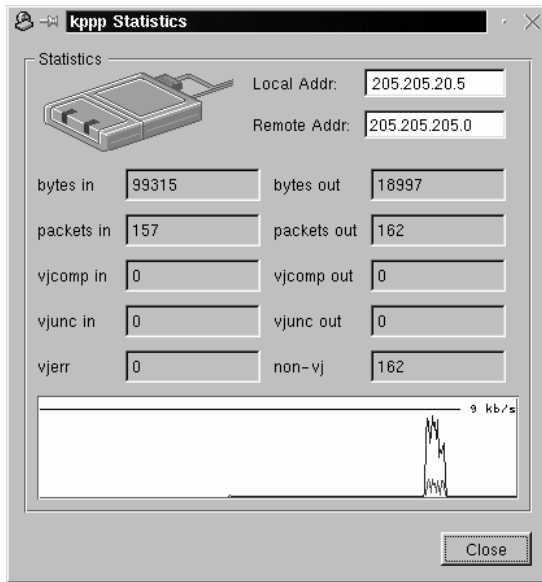
Figure 6–31 The Progress Window



Once you've established a connection with your ISP, Kppp will dock into the **Panel** or become minimized on the **Taskbar**, depending on the options you selected when you configured the application.


To view statistics about your connection, such as the throughput rate and IP addresses, open the status window of Kppp and click on the **Details** button, which will open a window showing you statistics about your connection (see Figure 6–32, *Connection Statistics*).

Figure 6–32 Connection Statistics



To disconnect, if the application is docked into the **Panel**, click once on the icon to open the connection status window (see Figure 6–33, *Connection Status*), then click on the **Disconnect** button.

Figure 6–33 Connection Status



If you chose to minimize the application to the **Taskbar**, then maximize the application and click on the **Disconnect** button.

7 The Web, E-Mail and News in GNOME and KDE

Once you've configured your PPP account, you can view websites for information, entertainment, software and more.

You can also send and receive e-mail -- the most popular online activity. And, you can trade information on **newsgroups**, which are essentially international bulletin boards covering thousands of subjects.

Before you can use e-mail or read newsgroups, however, you'll have to configure your e-mail and news applications, or **clients**. You can find two of the most popular mail and news clients in Netscape Communicator, included with your copy of Red Hat Linux.

In this chapter, you'll learn how to customize Netscape Communicator's browser, e-mail and news applications. A **suite** of applications, Communicator contains:

- Navigator -- for finding information on the World Wide Web
- Messenger -- for sending and receiving e-mail and news
- Composer -- for creating and publishing your own Web pages

We'll concentrate on Navigator and Messenger in this chapter.

Lynx: A lean and mean alternative browser

You'll probably use Navigator for your Web browsing. But your Red Hat Linux system also includes a useful browser called Lynx. What is Lynx? It's a "general purpose" browser for **character-cell displays**; it will display the content of pages on the Web, but not the images. It's very fast, and can be an excellent alternative for you if you're interested in reading information, without having to wait for graphics.

To give Lynx a try, open an Xterm window and type `lynx` at the prompt. You'll find a list of options at the bottom of the screen. To go to a site, for example, use the `[G]` key and enter the website address or URL; to quit, press the `[Q]` key.

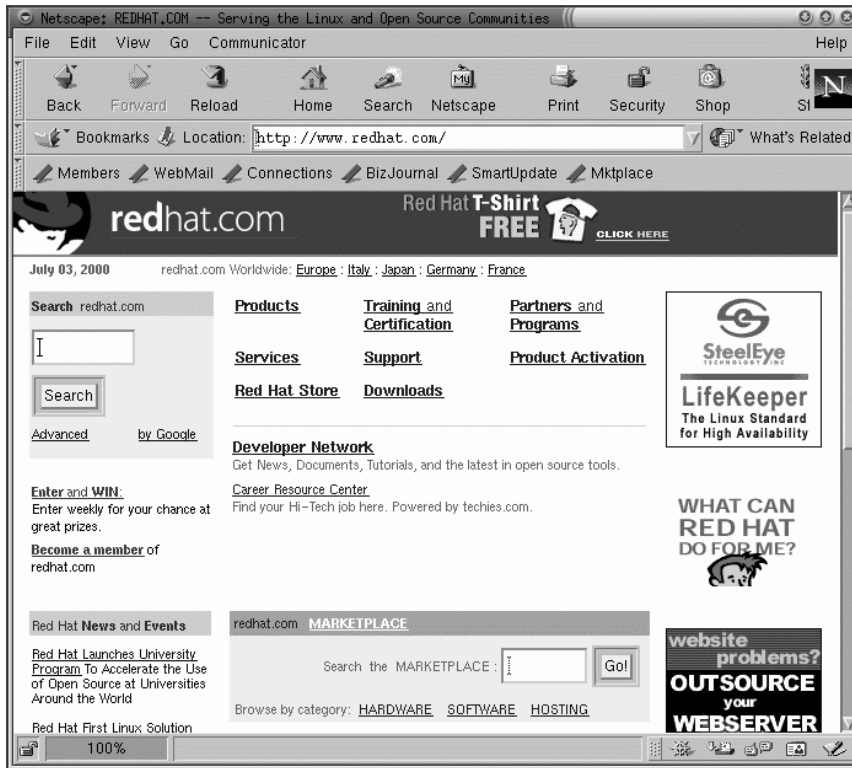
To learn more about Lynx, type `man lynx` at the shell prompt; or, if you want to print the document, type `man lynx | col -b | lpr`.

7.1 Using Navigator in GNOME and KDE

You'll find the Netscape Navigator **icon** on the **Panel** in GNOME and on the desktop in KDE.

To begin Navigator, click once on the icon. When the browser first starts, you'll be asked to read and agree to the software license. Click on the **Agree** button to proceed. Once the browser starts, it should look like Figure 7-1, *A Typical Navigator Window*.

Figure 7–1 A Typical Navigator Window



The top row containing the icons is called the **Navigation Toolbar**. Selecting the icons will perform the following actions:

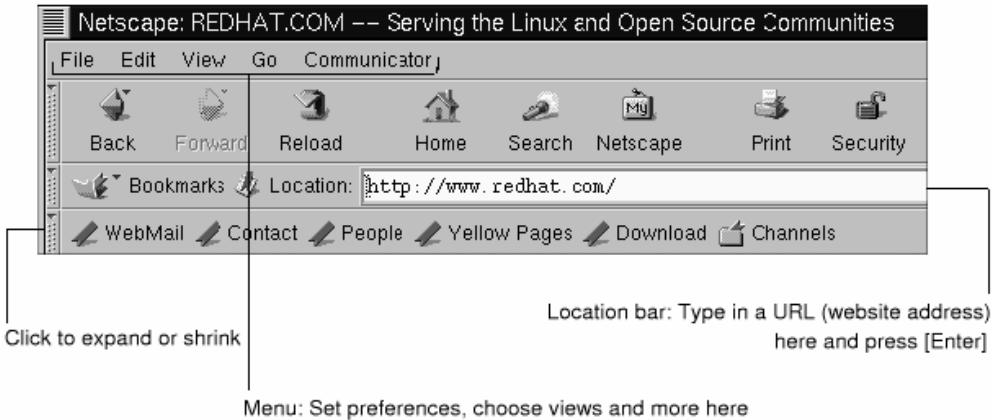
- **Back** -- Returns you to a previously viewed page. To jump back several pages, left-click on the icon and hold down the mouse button to reveal a pull-down menu of previously viewed pages. Drag and release the cursor to your chosen page.
- **Forward** -- When you've moved back to previously viewed pages, selecting this button will take you forward, toward your last viewed page.
- **Reload** -- "Refreshes" the current page in the browser. This can be useful, for example, if you want to view a page that's frequently updated.

- **Home** -- Selecting this button will immediately send the browser to the location specified as your **start page**. For more on how to set this feature to point to a different page, see Section 7.1.1, *Communicator Preferences*.
- **Search** -- Sends you to a **search engine** page. Search engines allow you to locate information, people, merchandise and more.
- **Print** -- If you have configured a printer, selecting this button will print the current page.
- **Security** -- Displays security information about the current page, as well as **encryption** information (which aids in security) about your browser.
- **Stop** -- Stops loading the page into the browser.

To browse other websites, just click on the **hyperlinks**, the differently colored, and usually underlined portions of a document you view.

You can also visit websites by typing in the **URL**, or website address, into the space provided in the **Location Toolbar** (see Figure 7-2, *Menu and Location Options in Navigator*). Click on the arrow pointing down at the right of the **Location Toolbar** to see a drop-down list of previously visited sites -- called the **history list**. To return to one of these sites quickly, select it from the list.

Figure 7–2 Menu and Location Options in Navigator



Opening locations in a new window

If you want to visit other sites, without losing your place on the page you're viewing, you can open a link in a **new browser window** by right-clicking on the link and selecting **Open Link in New Window** from the menu. (If you have a three-button mouse, or you selected three-button emulation during the installation of Red Hat Linux, you can also open a link in a new window just by clicking on the link with the middle mouse button.)

The following table provides keyboard shortcuts you can use in Navigator.

Table 7–1 Keyboard Shortcuts

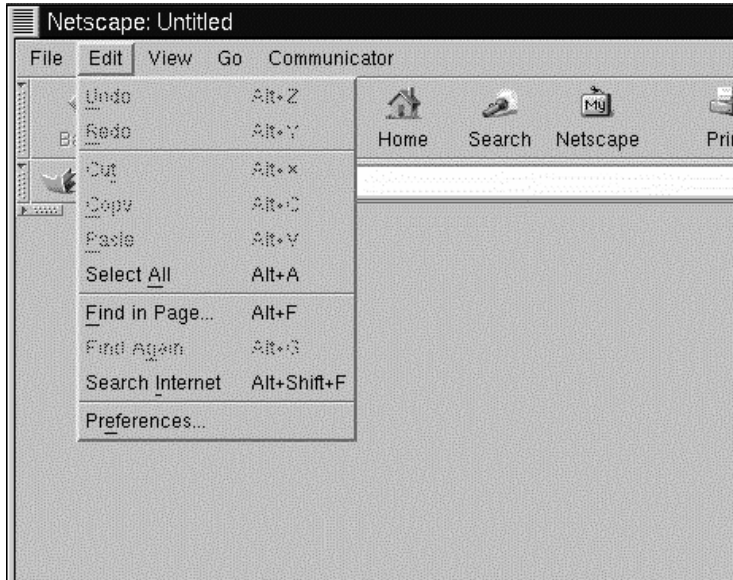
Command	Shortcut	Command	Shortcut
Stop Page Loading	[Esc]	Undo	[Alt]-[Z]
Reload	[Alt]-[R]	Print	[Alt]-[P]
Open Location	[Alt]-[L]	Open Page	[Alt]-[O]
New Mail Message	[Alt]-[M]	Edit Bookmarks	[Alt]-[B]
Find (on page)	[Alt]-[F]	Find (again)	[Alt]-[G]
New Window	[Alt]-[N]	Close Window	[Alt]-[W]
Quit	[Alt]-[Q]	Save As	[Alt]-[S]
Jump Forward	[Alt]-[Right Arrow]	Jump Back	[Alt]-[Left Arrow]

7.1.1 Communicator Preferences

Communicator allows you to customize many of its features. You can select how a Web page appears, security levels, and more.

To enhance your browsing experience, you can customize Navigator. From the menu, select **Edit=>Preferences** to open the options menu (see Figure 7–3, *The Preferences Item in the Edit Menu*).

Figure 7-3 The Preferences Item in the Edit Menu



The **Preferences** section is divided into subsections in the **Category** menu, including:

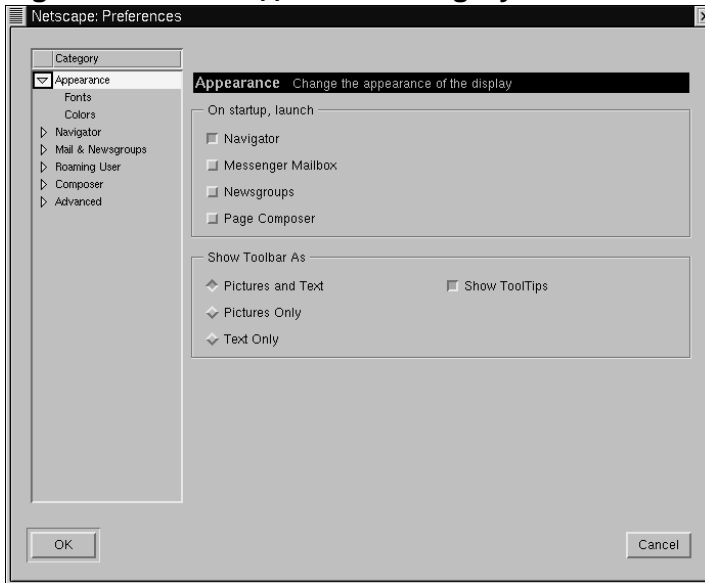
- **Appearance** -- Allows you to modify the way Navigator physically appears, such as the colors or use of icons on the navigation bars.
- **Navigator** -- Among other things, lets you specify your **Home page**, or the page you see when the browser starts. Also allows you to modify how long to keep previously visited sites in a **history**, for quick viewing later.
- **Mail & Newsgroups** -- Settings for using Communicator's e-mail and news clients. To learn how to use these settings, turn to Section 7.2, *E-Mail and News Settings for Messenger*.
- **Advanced** -- Allows you to specify settings for acceptance of **cookies**, which are small pieces of information stored on your computer and used to identify you to the websites you visit. Also allows you to modify other settings.

Each of these sections in the **Preferences** area can be expanded by clicking on the arrow next to the item on the menu. Let's look at each section in more detail:

The Appearance Setting

Click on the arrow next to the **Appearance** category to reveal all the related menus (see Figure 7-4, *The Appearance Category*).

Figure 7-4 The Appearance Category



In this category of the menu, choose among the following by selecting the box next to its listing:

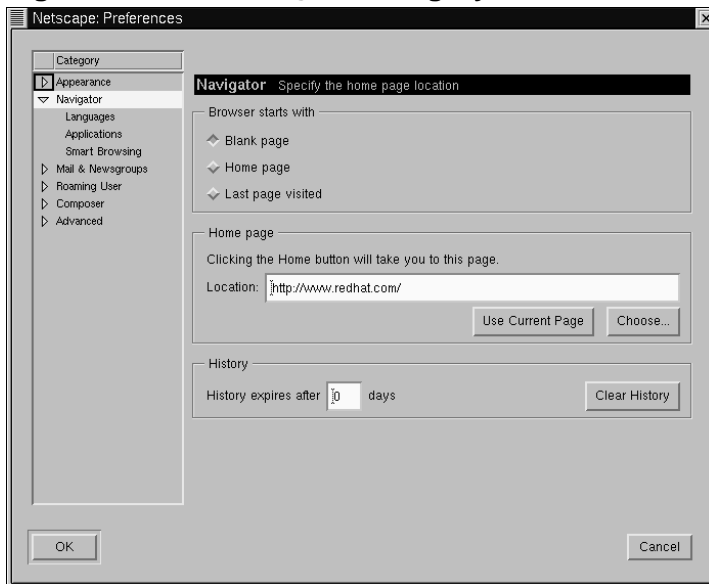
- **Appearance** -- Along with starting the browser, you can launch the e-mail application and newsreader clients.
- **Fonts** -- In this section, you can decide which **font**, or typeface, you want to use by default. If you're having difficulty reading the type on sites you visit, you can adjust the typeface and size of type here. Select your fonts and their sizes from the drop-down lists; your selections will then be reflected in the page you're currently viewing, so you can see their effect.

- **Colors** -- Lets you specify which colors you prefer for the text on pages, background and links to other sites. You can also choose whether you want links to be underlined in a document (links will appear only as colored text if you deselect this option).

The Navigator Category

In the **Navigator** category you'll find options for starting the browser, languages and related applications (see Figure 7-5, *The Navigator Category*).

Figure 7-5 The Navigator Category



On the first screen of the **Navigator** preferences, for example, you can choose to have the browser open with a blank page, rather than automatically go to a site when it starts.

In the **Home page** section of this category, you can specify a page you want to call the home, or start page. You can also specify that you want to have the browser automatically open the page you're currently viewing, by clicking on the **Use Current Page** button.

In the **History** section, you decide how long you want to retain previously visited sites in your history list. The history list can be a useful means of jumping to important sites you may have visited days earlier but neglected to bookmark. In the **History expires after** field, enter the number of days you want to retain URLs in your history file. To clear the file immediately, select the **Clear History** button.

Using your history file

To open your history file, from the menu go to **Communicator**
=> **Tools** => **History**.

The following category, **Mail & Newsgroups**, allows you to enter your account information to use Communicator's e-mail and newsreader clients. For more on these settings, turn to Section 7.2, *E-Mail and News Settings for Messenger*.

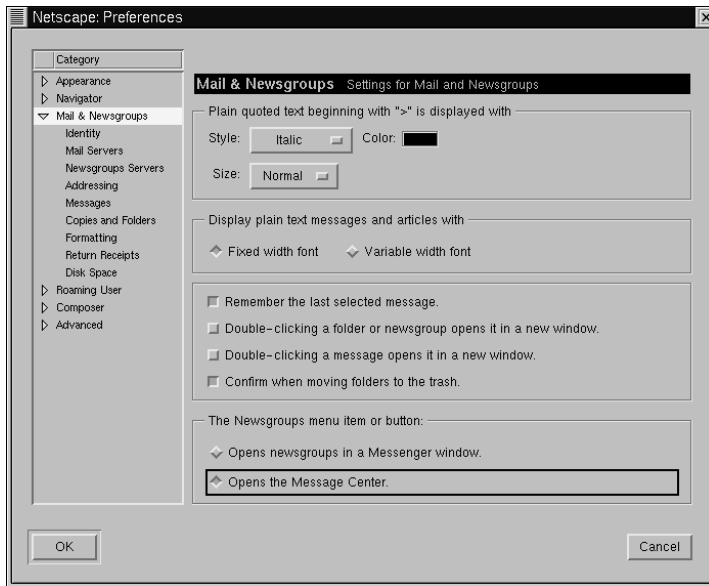
7.2 E-Mail and News Settings for Messenger

To read and send messages, you'll have to enter your settings in the **Mail & Newsgroups** category of Communicator (**Edit=>Preferences**). You'll also be able to save other related options in this category.

7.2.1 Entering Account Information

When you select **Mail & Newsgroups**, the category will expand, revealing sub-entries for your account information (see Figure 7-6, *The Mail & Newsgroups Category*).

Figure 7-6 The Mail & Newsgroups Category



General Options

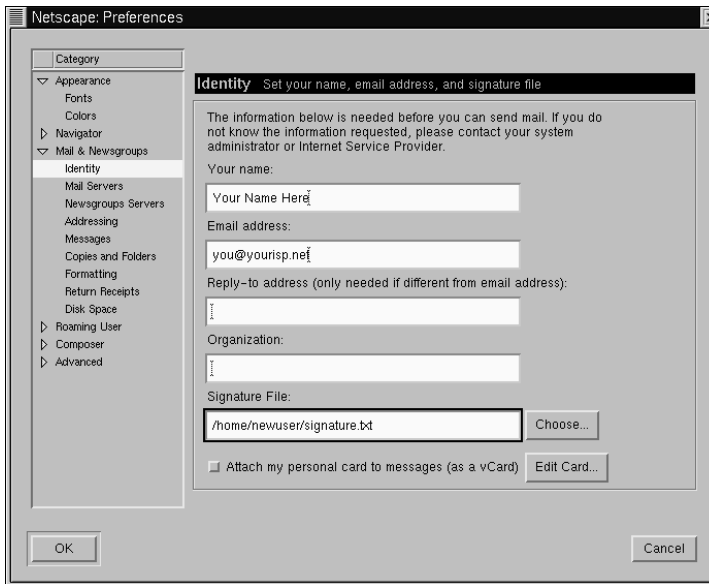
On the first dialog to the right of the categories, you can select how you want to have **quoted** messages appear. "Quoting" refers to material from a previous message, which appears as part of a follow-up message, such as a reply.

You can also select whether you want to view mail and news with a **fixed width** or **variable width** font. Selecting **Fixed width font** will preserve the display of fonts throughout messages and articles.

Account Identity

In the **Identity** category, enter your name, your e-mail address, the name of your organization (if applicable) and the path to your **signature** file (see Figure 7-7, *The Identity Category*).

Figure 7-7 The Identity Category



Your e-mail address

Often, your e-mail address is your login name for your account, followed by the @ symbol, then your ISP's name (as in "earthlink.net"). If you have questions about your address, you should check with your ISP.

If you want to direct users to reply to an e-mail address which is different from the primary address you're entering in this category, enter that different address in the **Reply-to address** field. When users reply, mail will be sent to this separate address. (This feature is optional, however; if you don't want this feature, you don't have to enter anything in this field.)

Creating a Signature

A signature is a brief note at the end of an e-mail or newsgroup posting. Often witty, signatures say something about the author of the message.

To create a signature, do the following, using the text editor Pico:

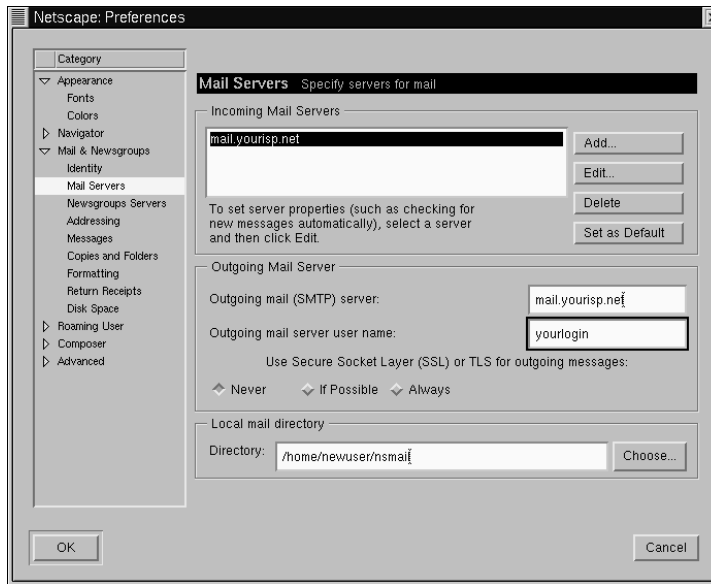
- Open an Xterm window.
- At the prompt, type `pico signature.txt`.
- In the window which opens, type the text you want for your signature (for example, **this is a signature from a new user**). You can write on more than one line, however **netiquette** frowns upon signatures longer than a line or two.¹
- Save the file by using the [Ctrl]-[X] keys. When you're asked whether you want to save the file, select the [Y] key for "yes."

Now, you'll have a new file, called `signature.txt`, in your directory. The full path to the file will be, for example, `/home/newuser/signature.txt`. When you write e-mail or post a message on a newsgroup, the text in `signature.txt` will automatically appear at the bottom of your message.

Mail Servers

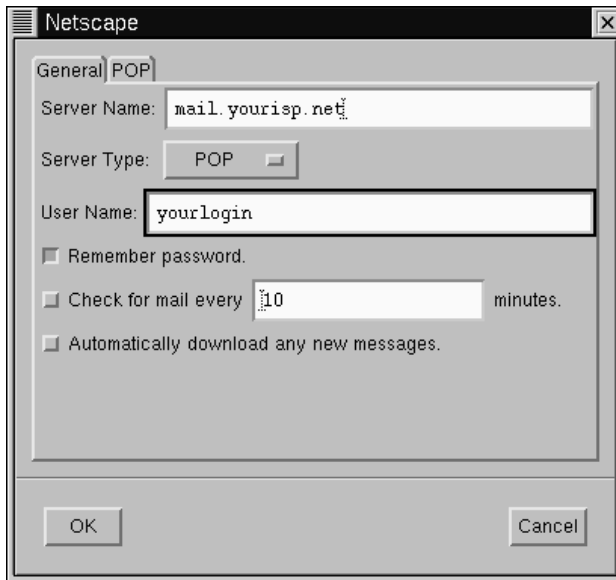
In the **Mail Servers** category, enter the server name through which your mail is sent and delivered (see Figure 7-8, *The Mail Servers category*).

¹ Netiquette means *Network Etiquette*, conventions of politeness which are recognized on the Internet.

Figure 7–8 The Mail Servers category

To add your account information, select the **Add** button. A new dialog, similar to Figure 7–9, *Adding a Mail Server*, will open.

Figure 7–9 Adding a Mail Server



In the **General** tab, enter:

- **Server Name** -- Where your mail is sent and retrieved from. Your ISP will be able to provide you with this information.
- **Server Type** -- From this drop-down box, choose **POP**, **MoveMail** and **IMAP**. Note that not all ISPs support all server types. Check with your ISP.
- **User Name** -- Enter your account login name.
- **Remember password** -- Check this option if you don't want to be required to enter your password each time you check your mail.
- **Check for mail every _____ minutes** -- Check this option and fill in the field to have Communicator check your ISP for new mail at specified intervals.
- **Automatically download any new messages** -- When you check this option, Communicator will automatically download new mail from your ISP's mail server and place it in your **Inbox**.

In the **POP** section:

- **Leave messages on server** -- Check this option if you don't want your messages to be deleted from your mail server when you retrieve them. Messages will remain on your machine, and will not be deleted on the server machine. Note that this option is unchecked by default; check with your ISP before deciding you want to store mail on their server indefinitely.
- **When deleting a message locally, remove it from the server** -- This option is also unchecked by default. If you select this option, mail will only be deleted off the server when you delete it on your machine.

When you're finished entering your mail account information, click the **OK** button to return to the **Mail Servers** window.

In the **Outgoing Mail Server** section, enter the name of the mail server. Your ISP should supply you with this information.

What is SMTP?

SMTP stands for *Simple Mail Transfer Protocol*, a protocol designed for the transfer of e-mail messages.

The option to **Use Secure Socket Layer** refers to encryption technology. For more information about security issues, refer to **Netscape's** online help guide.


You can specify where you want your mail to be stored by either typing in the location in **Local mail directory**, or by clicking on the **Choose** button, and selecting a directory in the navigation window which opens. By default, mail will be stored in a directory in your user directory; if your name is *newuser*, for example, the path to the directory will be `/home/newuser/nsmail`.

Unless you have a specific reason to choose a different directory, however, you can allow the default selection.

Newsgroups Servers

In the **Newsgroups Servers** category, enter the server name through which you can retrieve and post newsgroup messages (see Figure 7–10, *The Newsgroups Servers category*).

Figure 7–10 The Newsgroups Servers category



The procedure for entering the newsgroups server is somewhat similar to entering account information for e-mail services. In this instance, you must enter the name of the server to which you'll connect for access to newsgroups. Your ISP should provide you with this information.

Figure 7–11 Adding a Newsgroups Server

To begin, click on the **Add** button. In the dialog which opens (see Figure 7–11, *Adding a Newsgroups Server*), enter the newsgroups server information supplied by your ISP.

The checkboxes for **Support encrypted connections** and **Always use name and password** provide additional security options.

The first option, for **SSL**, is similar to the reference to **Secure Socket Layer** in the mail configuration. Check with your ISP to determine if this option is necessary.

The second option will cause **Messenger** to request your user name and password before connecting to your newsgroups server. This can be a useful in preventing unauthorized users from accessing newsgroups through your account.

Additional Options

The remainder of the categories in this section allow you to modify or accept settings for such options as message formatting, addressing, and disk space use.

To HTML or not to HTML?

Among the options you have for sending and receiving messages, you'll find **Message formatting**, in the **Formatting** category. This category allows you to decide whether you want to compose and send mail in **HTML** format.

HTML, or **Hypertext Markup Language**, is the format in which Web pages and some other documents can be viewed. It can provide some rich formatting tools for your messages -- for example, you can change the color or style of your fonts. However, not all users read their mail or news with applications that can adequately display HTML; netiquette also frowns on the use of HTML in mail and news.

Before sending your messages, make sure the recipients can read them. You can select the option **Ask me what to do...** in the section **When sending HTML messages...**, which will allow you to convert HTML messages to **plain text** for easier reading by everyone.

Once you've finished entering your settings and options in **Mail & Newsgroups**, you're ready to send and receive messages. Select the **OK** button to save your settings and exit the **Preferences** dialog.

Now, turn to Section 7.3, *Reading and Composing Mail in Messenger* to learn more.

7.3 Reading and Composing Mail in Messenger

Once you have configured your accounts in **Communicator**, you can begin using **Messenger** to send and receive e-mail and newsgroup postings (turn to Section 7.4, *Using Messenger to Read News* to learn how to set up your news preferences).

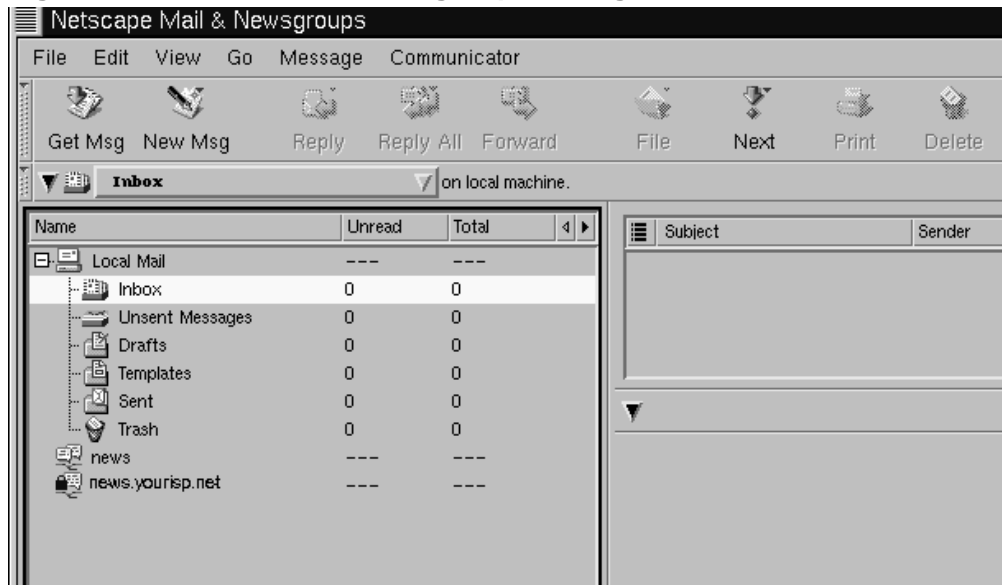
You can start **Messenger** in the following ways:

- Select **Communicator** => **Messenger** from the menu.
 - Press [Alt]-[2]
-

- From the **Component Bar** -- Select the mail icon, which looks like: 

The **Netscape Mail & Newsgroups** dialog which opens will look like Figure 7–12, *The Mail & Newsgroups Dialog*.

Figure 7–12 The Mail & Newsgroups Dialog



To compose e-mail, go to the **Toolbar** and click on the **New Msg** icon: 

You can also create e-mail by right-clicking on the **Inbox** and selecting **New Message** from the menu. Or you can use your [Alt]-[M] keys to quickly open the **Compose** dialog for e-mail.

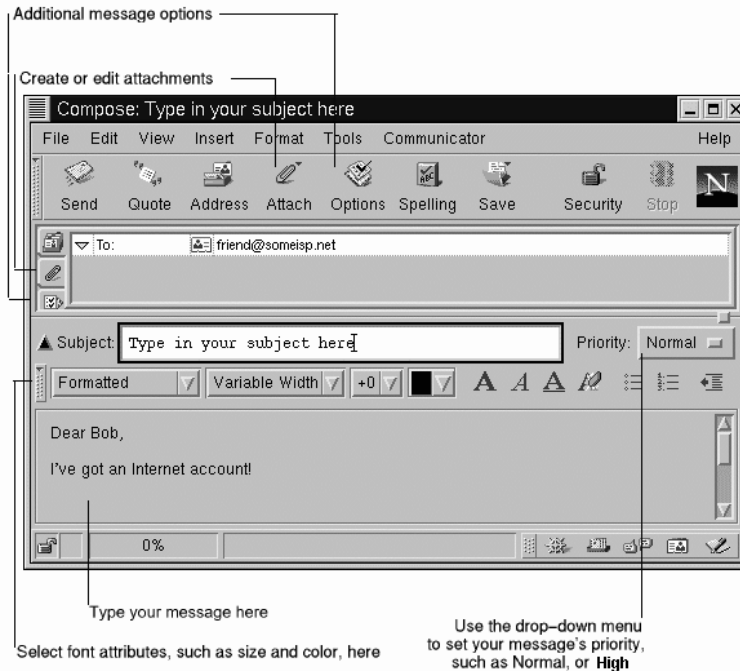
Creating a plain text letter

If you set your preferences to create e-mail and newsgroup documents in HTML format, press the [Shift] => **New Msg** icon to create a message in plain text.

Alternately, you can use the [Shift] => **New Msg** icon to create HTML-formatted documents if you selected to create your messages in plain text formatting.

The **Compose** dialog will look like Figure 7-13, *The Compose Dialog for E-Mail*.

Figure 7–13 The Compose Dialog for E-Mail



Note that the **title bar**, at the top of the **Compose** dialog, will show the subject of your letter (as in **Compose: Type in your subject here** in Figure 7–13, *The Compose Dialog for E-Mail*) once you save your message by clicking on the **Save** icon on the **Toolbar**.

In the address field, you can specify the following options:

- **To:** The primary recipient(s) of your letter
- **Cc:** To send a carbon copy to one or more recipients
- **Bcc:** To send blind carbon copies to recipients whose identity you want to remain private

- **Reply-To:** For specifying an e-mail address to which you want replies sent
- **Newsgroup:** To send your message to a newsgroup (such as `linux.redhat.misc`)
- **Followup-To:** To specify an e-mail address you want to use for any posts to newsgroups

In the address field, enter the e-mail address of your intended recipient, then press [Enter] to go to a blank address field. In each address field, you can specify address options by selecting one from the drop-down menu next to the address.

When you're finished typing in your addresses, press the [Tab] key to go to the subject field. Type the subject of your letter in this field and press the [Tab] key.

Now, type the body of your letter. When you're finished, click on the **Send** icon on the **Toolbar** to send your e-mail.

That's it for creating and sending basic e-mail. However, **Messenger** provides some useful extras for **attaching** files, making sure that your letter has been received and more.

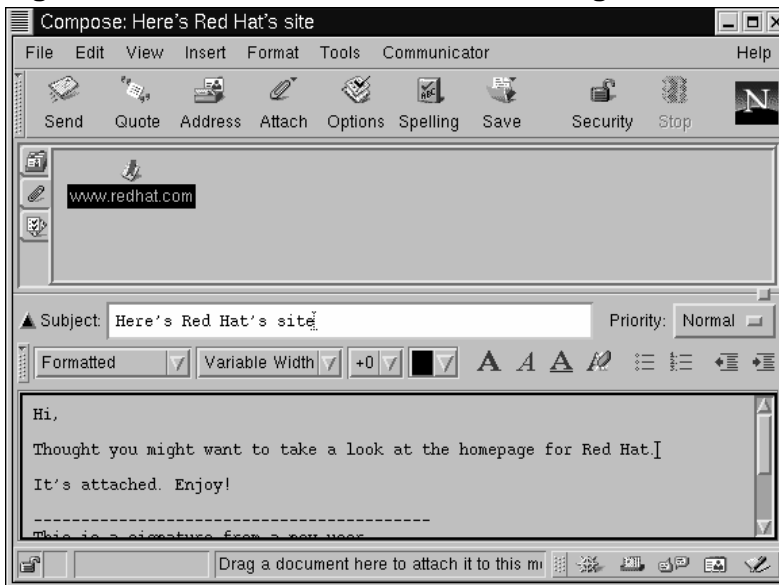
7.3.1 Other E-Mail Options

If you want to send someone a Web page or a file on your computer, you can attach the document to your message. Click on the paperclip icon either on the **Toolbar** or on the tab beneath the address field.

For example, if you want to attach a page from the Web, click on the paperclip icon, and select **Web Page** from the menu. In the dialog that opens, type the **URL**, or website address, as in Figure 7-14, *The Attach Web Page Dialog*.

Figure 7–14 The Attach Web Page Dialog

When you have finished typing in the URL, click on the **Attach** button. In your letter, you'll see that the URL has been included with your letter (see Figure 7–15, *E-Mail and Attached Web Page*).

Figure 7–15 E-Mail and Attached Web Page

When you send your letter, the page associated with the URL you had typed will be sent with your letter, as well.

Drag and drop Web pages


You can also "drag" a URL from Navigator's **Location Toolbar** and "drop" it in your letter. In your letter, click on the paperclip icon under the address field. Now, from Navigator, go to the **Location Toolbar** and hold down your mouse button over the nearby bookmark icon, which looks like this:



Now, drag the icon to the attachments field of your letter.

In the same way that you attach Web pages to your letter, you can attach files on your computer. When you select **Attach** from the **Toolbar** or beneath the address field, choose **File** from the menu. The **Attach File** dialog will open, showing you the contents of your login directory (see Figure 7-16, *The Attach File Dialog*).

Figure 7-16 The Attach File Dialog



Drag and drop files

You can drag files from your desktop or from your file manager in the same way that you can drag URLs from Navigator. Just hold your mouse button over the file you want to attach, and drag it to the attachments field of your letter.

Have you ever wondered whether your e-mail was actually read, or even received? Messenger can send you a **return receipt** to confirm that your mail arrived and was opened.²

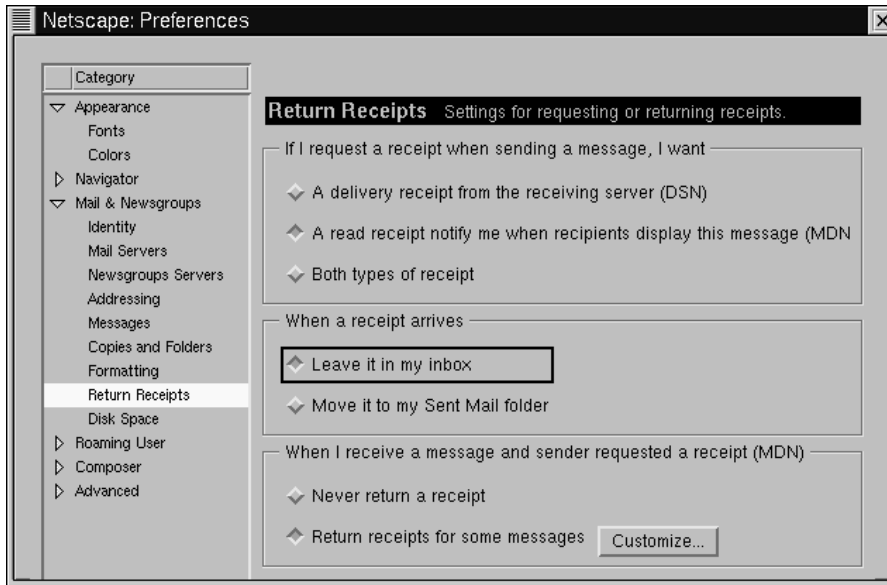
Messenger supports two kinds of return receipts:

- Delivery receipts -- notify you that your mail was delivered to the inbox of your intended recipient on their mail server
- Read receipts -- notify you that the recipient has opened your message

To specify preferences for return receipts, go to **Edit => Preferences => Mail & News-groups => Return Receipts**. In the dialog which opens (see Figure 7-17, *The Return Receipts Dialog*), choose how you want to use return receipts for mail you send and mail you receive.

² For return receipts, the recipient's server must support the **Delivery Service Notification (DSN)** standard. For read receipts, the recipient's server must support the **Message Disposition Notification (MDN)** standard. So, you may not be able to receive return receipts, depending on your recipient's e-mail client and its configuration.

Figure 7–17 The Return Receipts Dialog



To activate receipt notification, in your message go to **View => Options** and check **Request return receipt**. You can also click on the options tab beneath the attachments tab, or click on the **Options** menu on the **Toolbar**.

For more information about e-mail options, such as address books, filtering e-mail and more, refer to Netscape Communicator's documentation, in the **Help** menu.


7.4 Using Messenger to Read News

One of the best methods to distribute information to a vast audience is through **Usenet**. Comprised of thousands of **newsgroups**, Usenet is used to exchange information, gossip, and more.³

³ If you haven't used Usenet before, you should be aware that newsgroups can contain information and material that might be offensive to you. Be cautious about which newsgroups you subscribe to.

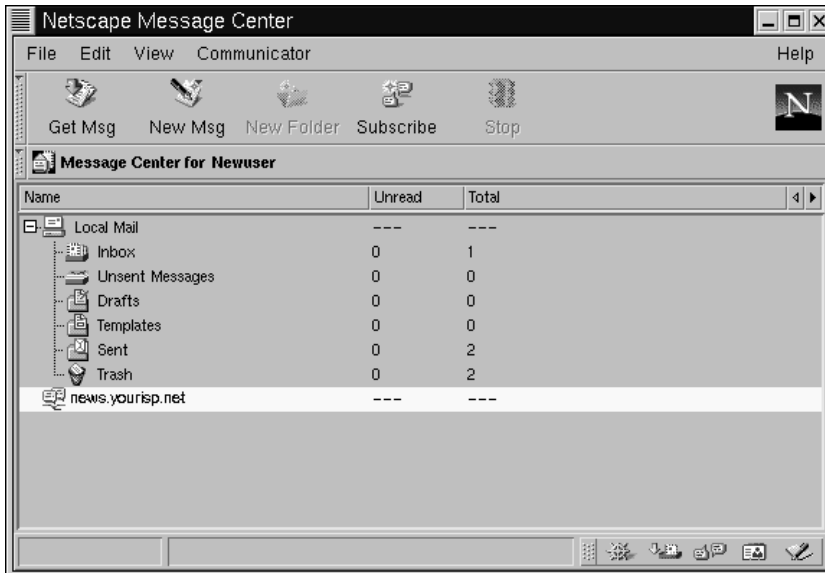
To read newsgroups, you'll need a **news reader**. One such news reader is **Messenger**, which lets you read and post to newsgroups similar to the way you work with e-mail.

You can start Messenger in the following ways:

- Select **Communicator** => **Messenger** from the menu.
- Press [Alt]-[2]
- From the **Component Bar** -- Select the newsgroups icon, which looks like: 

The **Netscape Mail & Newsgroups** dialog which opens will look like Figure 7-18, *The News Server Entry in Messenger's Message Center*.

Figure 7-18 The News Server Entry in Messenger's Message Center



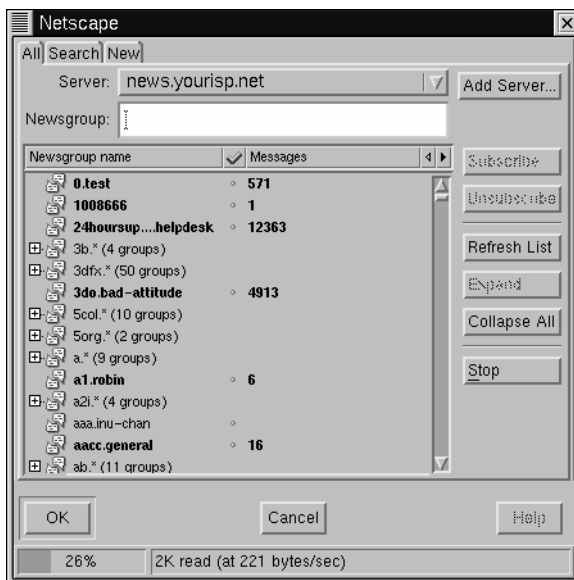
If you haven't configured a news server

If you haven't already configured your news server entry, turn to *Newsgroups Servers* in Section 7.2.1.

Your next task is to download the list of newsgroups from your ISP's news server, and to **subscribe** to those groups in which you have an interest. Be prepared: Depending on the speed of your connection, downloading the list of many thousands of newsgroups can take some time. While you wait for the complete list, you may have enough time to make a pot of coffee, phone a friend -- or read a novel (just kidding).

To begin, click on **Subscribe** on Messenger's menu, or right-click on the news server entry in the **Message Center** dialog, and select **Subscribe to Newsgroups**.


Figure 7–19 The Newsgroups Subscription Dialog



In the tabbed window which opens (see Figure 7–19, *The Newsgroups Subscription Dialog*), you'll be able to:

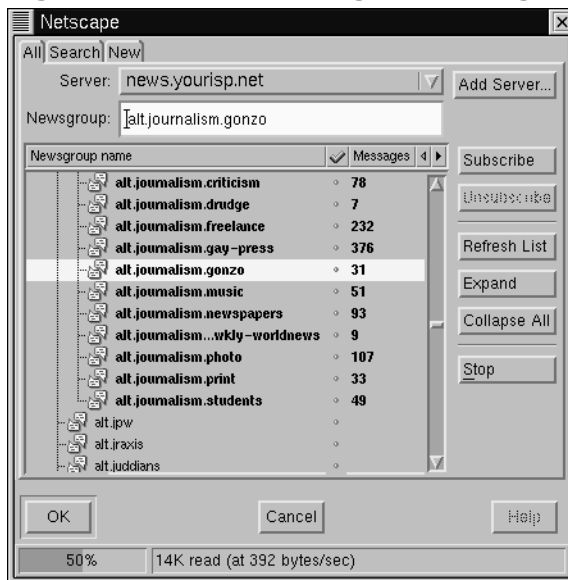
- See all of the available newsgroups to which you can subscribe (**All**)
- Search for a particular newsgroup by full or partial name (**Search**)
- View newsgroups which have been added by your ISP since you last downloaded (or **refreshed**) the full list of newsgroups (**New**)

An indicator on the bottom of the dialog will show your progress as you download the complete list of newsgroups.

Groups shown in bold have new messages. Listings which are preceded with a "+" are the top level of groups related to that particular subject, for example:  alt.*.journalism.* (50 groups)

To see these related groups, click on the "+" to expand the list.

Figure 7–20 Subscribing to a Newsgroup



Click to highlight the name of a group that interests you, then click on the **Subscribe** button (as in Figure 7–20, *Subscribing to a Newsgroup*).

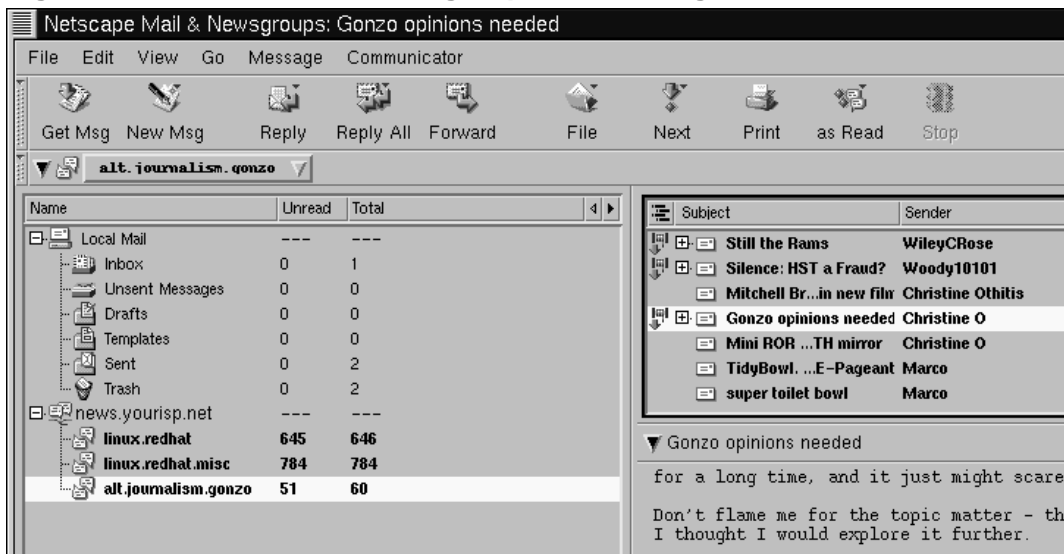
You can also search for newsgroups based on your interests in the **Search** dialog. Let's say you're an avid news and media fan. In the **Search For** field, type **journalism**

and click on the **Search Now** button (or press [Enter]). Your search will likely result in quite a few newsgroups to which you can subscribe.

How to unsubscribe

You can unsubscribe to newsgroups in the same way that you subscribe to them.

Figure 7–21 Subscribed Newsgroups in Messenger



When you are finished subscribing to newsgroups, click the **OK** button. Your selections will appear in **Messenger's** main dialog (see Figure 7–21, *Subscribed Newsgroups in Messenger*). Note that the **Total** category represents the total number of messages in the group; the **Unread** category represents the number of messages you have not read.

To view a newsgroup, double-click on its name. In the top right pane of the dialog, you'll see the messages' subjects; in the bottom right pane, you'll be able to read the associated messages.

To write a message to your selected newsgroup, click on **New Msg** on the menu, or go to **File => New => New Message**. Like writing e-mail, enter your message in the **Compose** dialog which opens and send your message. It will be **posted** to your current newsgroup.

Either by selecting one of the menu items or by right-clicking in the message body, you can perform the following on posted messages:

- **Reply** -- Send a reply to the poster either in the newsgroup or privately, by e-mail
- **Reply All** -- Reply to both the poster and the newsgroup
- **Forward** -- Forward the message to an e-mail address
- **File** -- Place the message in a **Messenger** folder, such as your **Inbox**
- **Next** -- Move to the next message on the list
- **Print** -- Print the message
- **as Read** -- Mark the message as **read**, in which case it would no longer be highlighted as a new, or unread, message

There are other applications which can be useful and fun ways to chat and exchange information online. For example, you may wish to explore applications such as messengers and **IRC** clients (IRC means **Internet Relay Chat**). But the staples of browsing (e-mail and news reading) may keep you entertained and busy for quite some time.

8 Customizing Your Desktop in GNOME and KDE

Both GNOME and KDE offer you plenty of ways to add your own personality to your workspace. Each environment provides powerful **control centers**, which let you to change the way your workspace looks, sounds, and more. In many respects, each control center behaves identically.

Don't like the color of your desktop? Change it to a different color -- or add a background photo. Want something more expressive? Change your **theme**, the coordinated look and "feel" of your workspace.

Both control centers offer quite a few options for you to modify. For more detailed instructions on using each, you should click on the **Help** buttons in each or read the related documentation from the environments' help browsers.

8.1 Customizing with the GNOME Control Center

You can start the GNOME Control Center from the **Main Menu Button (Programs => Settings => Gnome Control Center)**.

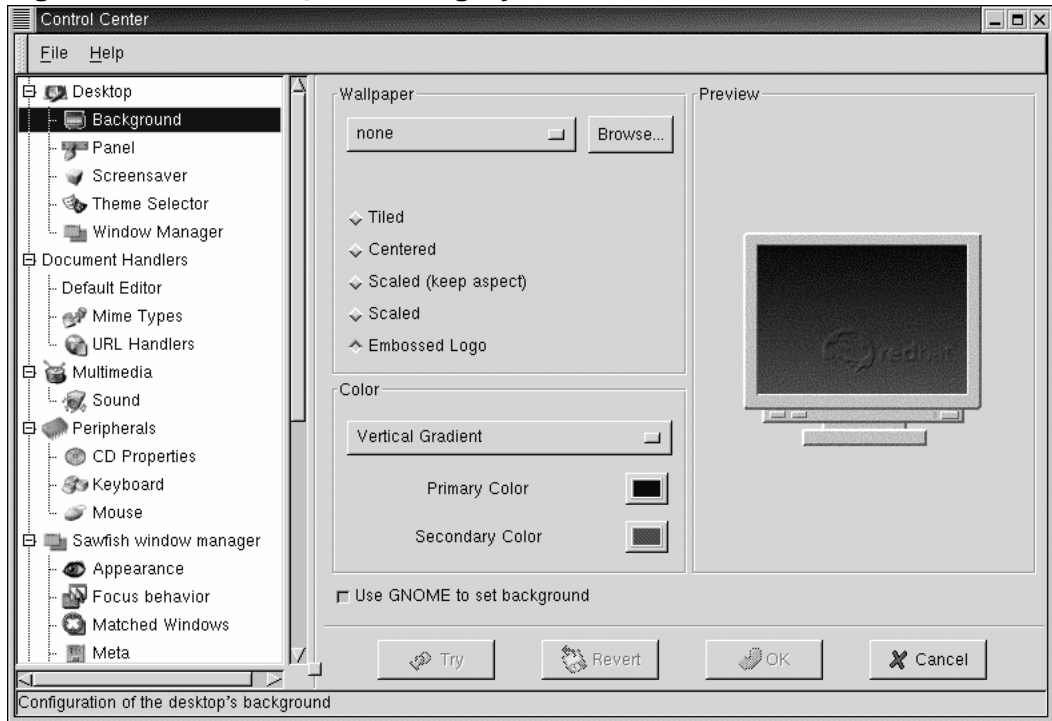
You can also start the GNOME Control Center from the **Panel**, by clicking on the

icon, which looks like:



8.1.1 Changing the Background

When the GNOME Control Center first starts, it will look like Figure 8–1, *The Background Category in the GNOME Control Center*.

Figure 8–1 The Background Category in the GNOME Control Center

The GNOME Control Center is divided into two panels: On the left, you'll find categories, and on the right, in the main panel, the corresponding elements you can modify for a category. The categories may be **collapsed** when you first open the control center. Click on the **+** to **expand** the categories.

Try it on for size

Not sure you'll be happy with your selections? Click on the **Try** button to see what your desktop will look like with your changes. If you don't like it, just click on the **Revert** button to return to your original settings.

To change the background of your desktop, click on the **Background** category. You'll see your current desktop background in the preview screen of the main panel.

Select the color button under **Primary Color** to change a solid color of your desktop. If you selected either **Gradient** option, when you're finished adjusting the first color, click on the corresponding color box under **Secondary Color** to adjust the second color. You can also decide whether you want to gradient to be oriented horizontally or vertically on your screen.

If you want something a little more expressive, try adding **wallpaper** to your desktop. Rather than a one- or two-color background, wallpaper places an image on your desktop, which can be:

- **Tiled** -- The image is repeated across the desktop
- **Centered** -- The image is centered on the desktop, surrounded by your chosen background color or colors
- **Scaled (keep aspect)** -- The image is **scaled** to fit as much of the screen as possible, without image distortion
- **Scaled** -- The image is stretched to fit the entire screen
- **Embossed Logo** -- Uses the default Red Hat logo

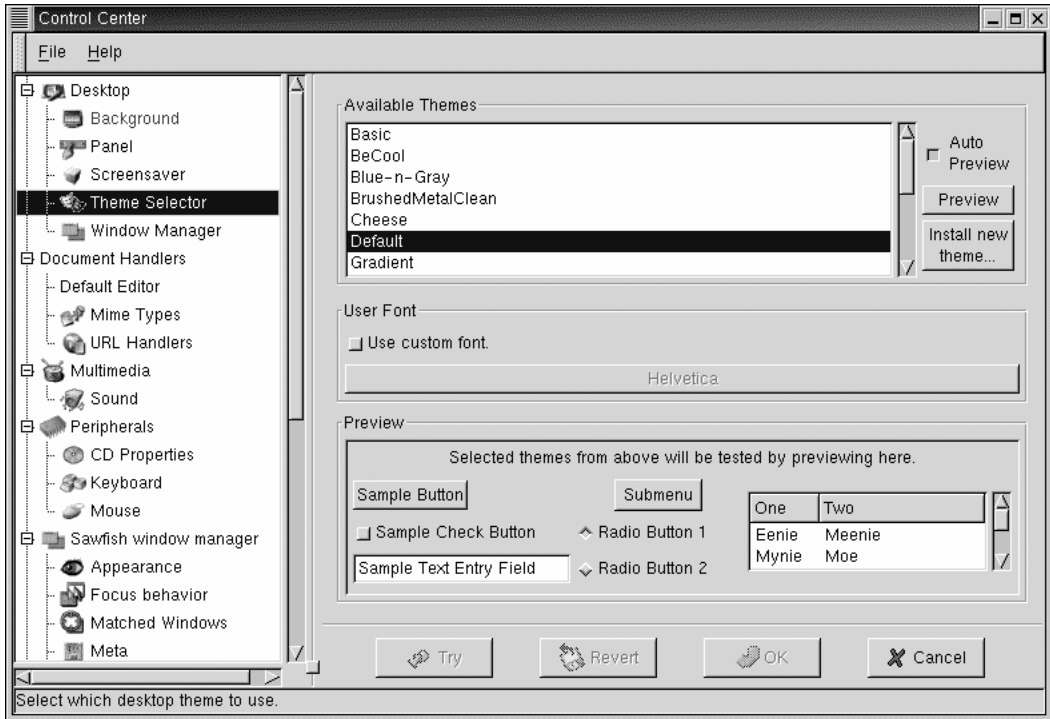
Finally, deselect the option to **Use GNOME to set background** if you would prefer to set the background through another means (such as through your window manager's configuration program).

8.1.2 Changing the Theme

You can change the look and "feel" of your workspace with themes. The **Panel**, text boxes, buttons and other visual elements will appear coordinated with the theme.

To change your current theme, click on the **Theme Selector** in the category panel.

Figure 8–2 The Theme Selector Category



The right panel is divided into sections which allow you to select a new theme, preview its effect and modify the default font that theme will use (see Figure 8–2, *The Theme Selector Category*).

If you select the **Auto Preview** option, you will be able to preview what the theme will look like in the **Preview** section. If the option is not selected, you can still preview the theme by clicking on the **Preview** button.

More themes available

Don't see a theme you like? You can find more themes at *gtk.themes.org*, <http://gtk.themes.org>, a repository for many themes. To **install a theme**, download a theme from the website, then select the **Install new theme** button. In the dialog which opens, scroll to the name of the file you downloaded (which will probably be in your login directory, named something like *New_theme.tar.gz*). Left-click on the file name, then click the **OK** button. Your new theme will appear in the **Available Themes** section.

In the **User Font** section, click on the option to **Use custom font**, then click on the activated font box to open a new dialog which will let you select a different font on such items as the menu and Tasklist.

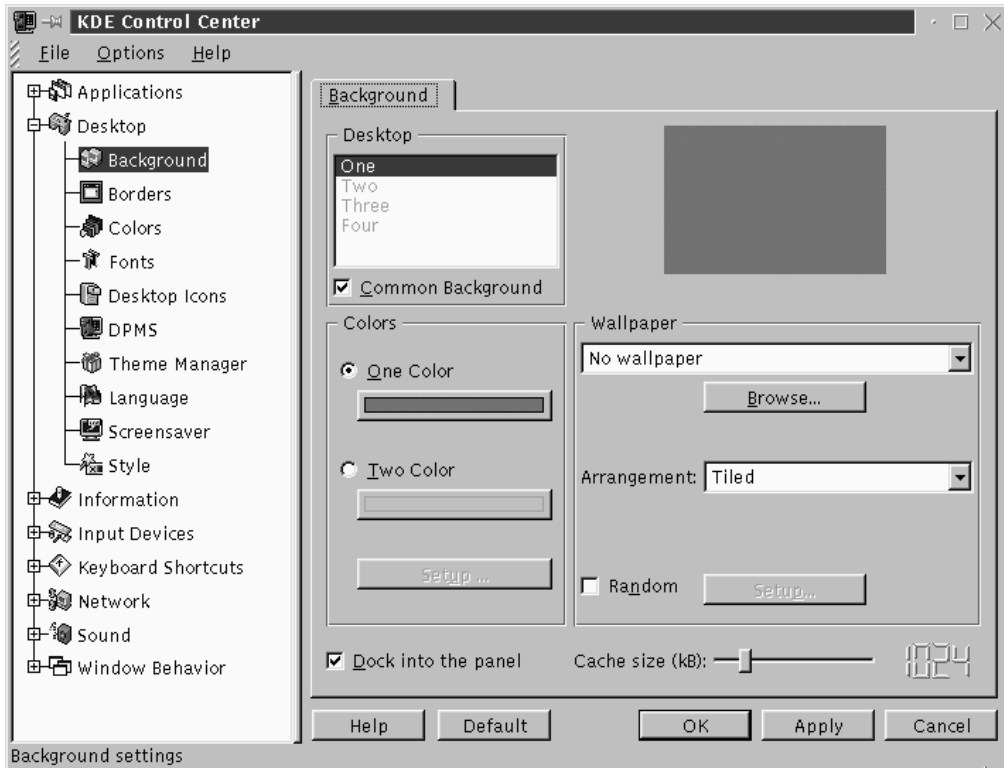
8.2 Customizing with the KDE Control Center

You can start the KDE Control Center from either the **Main Menu K** icon (**Settings => Desktop**) or from the **Panel**, by clicking on the KDE Control Center icon, which looks



8.2.1 Changing the Background

The **Background** category of the KDE Control Center will look like Figure 8-3, *The Background Category in the KDE Control Center*.

Figure 8–3 The Background Category in the KDE Control Center

The KDE Control Center is divided into two main panels: On the left, you'll find categories, and on the right, in the main panel, the corresponding elements you can modify for a category. The categories may be collapsed when you first open the control center. Click on the + to expand the categories.

In the **Background** category, you can change the background colors from a single color to a two-color gradient, and select whether the gradient is horizontal or vertical on your desktop.

Automatic preview

Whenever you make a selection for your background, you'll automatically see a preview of the effect in the dialog.

A unique feature of the **Background** category in the KDE Control Center is the option to create distinct backgrounds for each desktop. When the **Common Background** option is unchecked in the **Desktop** panel, the labels for each of your desktops will be highlighted. Click on a label to customize the corresponding desktop. (Checking the **Common Background** label will apply your background selections to all your desktops. Note that you can add both colors and/or **wallpaper** to each of your desktops, as long as the **Common Background** option is unchecked.

Click on the drop-down list in the **Wallpaper** section to choose a pre-selected pattern or photo for your desktop. In the drop-down list next to **Arrangement**, you can select how you want the image to appear: centered, **tiled** (the pattern is repeated across your desktop), and others.

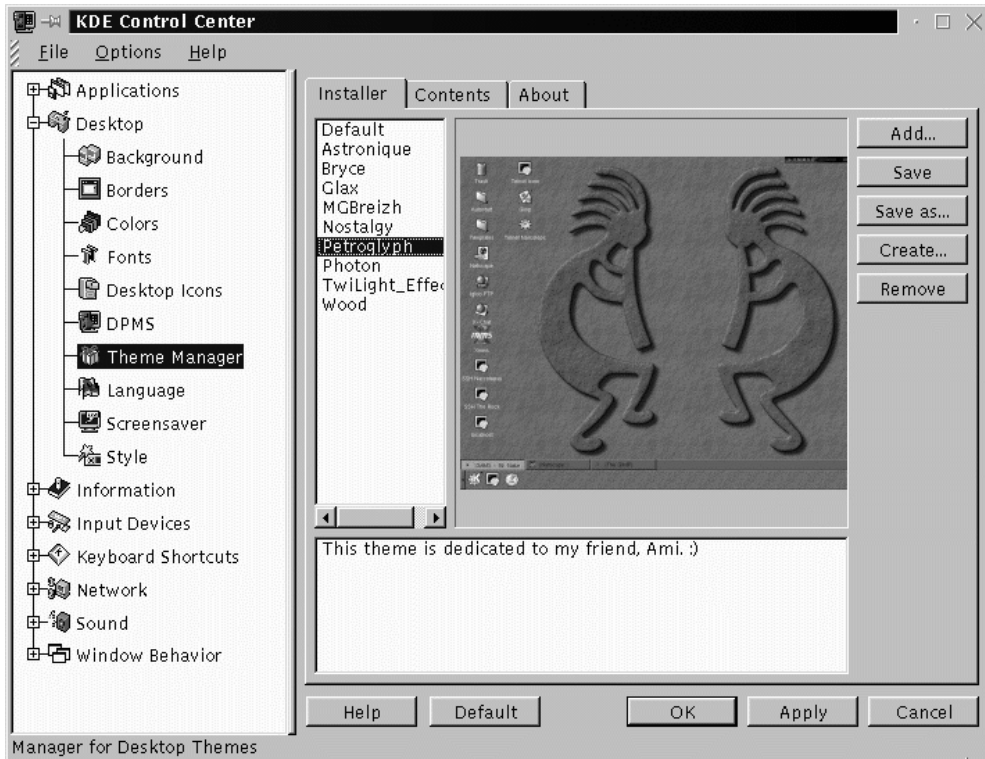
If you want something closer to a total makeover, try changing the theme of your workspace.

8.2.2 Changing the Theme

Applying a new theme to your workspace can change the look and feel of your **Panel**, buttons, menus, and even the background.

To begin, select the **Theme Manager** category (see Figure 8–4, *The Theme Manager Category*). By default, you have several themes, however there are quite a few more themes available at *kde.themes.org*, <http://kde.themes.org>.

Figure 8-4 The Theme Manager Category



If you want to install a theme you have downloaded, click on the **Add** button. In the dialog which opens, locate the name of the theme file you have downloaded then either double-click on it or click on it to highlight the file, and click the **OK** button. Your theme will be installed as one of the themes you can try.

9 A Short Introduction to the GIMP

The GNU Image Manipulation Program (GIMP) can help you create, alter, manipulate and enhance digital image files -- photographs, scanned images, computer-generated images, and more. This chapter is intended to give you a quick overview of the GIMP and its more common applications.

Like other fields, image manipulation has its own vocabulary. This chapter can't possibly provide complete documentation for all of the the terms and concepts needed for complete understanding of the GIMP. If you want to learn more, a variety of books are available which document the GIMP and its capabilities, as well as the background concepts of image manipulation. Please refer to Section 9.3, *Where to Find More Information* if you need more information about the GIMP.

As you'll quickly find out if you try the GIMP, it includes a wealth of tools for editing images -- much more than could be described in a single chapter.

Furthermore, if you can't create the right effect with the tools that the GIMP provides, and you're not afraid to write scripts in the Scheme scripting language, you can add a new **Script-Fu** to the GIMP. A Script-Fu is somewhat analogous to a macro in a word processing or spreadsheet application -- it applies a set series of effects (filters or other GIMP capabilities) to a graphic or piece of text. Anyone can use the pre-written Script-Fus included with the Gimp, but you can also write your own. See Section 9.3, *Where to Find More Information* for more sources of information about Script-Fus.

The GIMP can do a lot of things, but it has some limitations. Graphic artists who create images for print media may find the GIMP lacking, or just more difficult to use, for pre-press applications.

For example, the GIMP can't (yet) work in the CMYK (cyan, magenta, yellow and black) colorspace, which is generally used for printing color graphics. The GIMP instead uses the RGB (red, green and blue) colorspace, which should be familiar to people who have created graphics for the Web.

Since the GIMP is well-suited for creating graphics for the Web, that's what we're going to do in this chapter. First, we'll edit a photo for display on a webpage. Then

we'll create a background image, which will be tiled on the background of our webpage. Finally, we'll use the **GIMP** to create a title image for the webpage.

Before we start "learning by doing," you'll need to know some basics: creating files, saving files, the Toolbox and Dialogs.

9.1 GIMP Basics

To use the **GIMP**, you'll need to know some of the basics. From an Xterm, you start the **GIMP** using the command (you guessed it): `gimp`.

Like many applications, the **GIMP** provides more than one method to accomplish tasks. Generally, there are three ways to work in the **GIMP**:


- By right-clicking on an image, which will display a set of menus containing most of the **GIMP**'s many capabilities
- From the Toolbox, using the tools or the **File** or **Xtns** pull-down menus -- see Section 9.1.1, *The Toolbox*
- With keyboard shortcuts -- see Section 9.1.5, *Helpful Keyboard Shortcuts*

When you start the **GIMP**, you'll see the Toolbox.

9.1.1 The Toolbox

The **GIMP** Toolbox is a palette of commonly used **GIMP** tools:

Figure 9–1 The Toolbox



You can minimize the Toolbox (send it to the panel), but if you close the Toolbox, you're closing the GIMP.

If you move your mouse over the individual buttons in the Toolbox, a pop-up dialog will appear to identify the tool's purpose.

Click on the button for the tool you want to use in the Toolbox. The button will then be colored in, to indicate that it is the active tool.

Generally, you use a tool on an image by activating the tool, then clicking and dragging on your image. If a portion of the image is selected, then the tool will only work on the selected portion. Selections are identified by a moving dotted-line border (called "**marching ants**").

Many of the tools have options. Double-clicking on the tool's button will display the **Tool Options** dialog for that tool.

Tools in the Toolbox



Select rectangular regions

Used to select rectangular portions of an image. If you click on the image, then press and hold [Shift] while you're dragging, the rectangle will be a square. If you click on the image, then press and hold [Ctrl] while you're dragging, your starting area will be the center of the rectangle (normally the starting point is at an edge of the rectangle).



Select elliptical regions

Used to select circular or oval portions of an image. If you click on the image, then press and hold [Shift] while you're dragging, the ellipse will be a circle. If you click on the image, then press and hold [Ctrl] while you're dragging, your starting point will be the center of the ellipse (normally, your starting point is at an edge of the ellipse).



Select hand-drawn regions

This tool, also known as the **lasso**, is used to draw a freehand selection. Click and drag to draw a selection. When you get back to where you started, release the mouse button.



Select contiguous regions

This tool, also known as the **fuzzy selection tool** or **magic wand**, is used to select portions of an image based on color values. When you click once on an image with the magic wand active, pixels with color values close to that pixel will be selected. If you click and drag the magic wand, a larger group of pixels will be used as the starting point. All regions of the image with a color value similar to the selected group of pixels will be selected.



Select regions using Bezier curves

This selection tool is similar to the freehand select tool, but it allows you to customize your selection. If you need to be very precise when you're selecting

a complicated object, the Bezier tool will probably be the right choice for the job.

To create a Bezier selection, click to create control points in a connect-the-dots pattern around your selection. Put a control point at every corner or angle of your selection.

When you get around to where you started, click on the first control point you created to close the selection. Now you can move the control points and create curves between them to modify your selection.

To move control points, hold down the [Ctrl] key and drag the point you want to move.

To create curves, you need to click and drag on a control point. As you drag away from the control point, you'll see handles appear. Click and drag on the handles to control the shape of the curve.

If you only want to affect the curve on one side of a control point, hold down the [Shift] key and drag on one of the handles.

Once you've finished modifying the Bezier selection, click inside the selection to activate it.



Select shapes from image

The **Select shapes from image** tool (or **intelligent scissors**) is still in development. Once it is perfected, it is intended to improve upon your selection and automatically select the correct shape. This tool occasionally behaves unpredictably, but give it a try. Obviously, this type of selection is easier for the GIMP to handle if the area selected has a distinct edge.

Click on the intelligent scissors tool in the toolbox to activate it. Then, as if you were using the Bezier tool, click on the image to draw a connect-the-dots pattern around your intended selection. The intelligent scissors tool will try to guess what you're selecting, so the lines it draws between the control points will (hopefully) follow the edge of your selection instead of being just a straight line.

To refine your selection, you can move the control points by clicking on one and dragging. Once you're happy with the selection, click inside it to activate it.



Move layers and selections

The **move tool** is used to move a selected region.

Just after you've created an active selection, if you move your mouse cursor over it, you'll see the mouse cursor turn into the move symbol. Click and drag on the selection to move it. This technique works even if the move tool is not activated in the Toolbox.

With the move tool activated, if you click and drag on your image, you'll move the entire image (or the active layer). If you press the [Alt] key and then click and drag on the image, you'll just move the selection (the marching ants outline).

With the move tool activated, if you press [Alt] and then use the arrow keys, you can move the currently active selection, one pixel at a time.



Zoom in & out

Also known as the **magnifying glass**, the **zoom tool** is used to enlarge (by clicking) or reduce (by [Ctrl]-clicking) the image. You can also click and drag to zoom in on a specific region in the image.

Alternatively, the [=] keyboard shortcut will zoom in and [-] will zoom out.



Crop the image

The **crop tool**, which resembles a scalpel, is used to cut away portions of an image. Click and drag to draw a rectangular selection that you want to keep. The **Crop Information** dialog will appear. Click on the **Crop** button on the **Crop Information** dialog to cut away everything outside the selection. Alternatively, click on the **Selection** button to make a customizable crop box appear. The control squares at the edge can then be individually moved to refine the crop selection.



Transform the layer or selection

The **transform tool** can be used to rotate, scale, shear or change the perspective of an image or selection. Double-click on the **Transform** button to display the **Transform Tool Options** dialog. Choose which action you want to perform on the image or selection. Move the mouse cursor over the selection or image and you'll see the mouse cursor change into the transform tool circular arrows icon. Then click and drag on the image image or selection to perform the action that you've chosen.



Flip the layer or selection

The **flip tool** is used to flip (or produce a mirror image) of an image or selection.



Add text to the image

The **text tool** is used to add text to the image. Select the text tool and then click on the image. The **Text Tool** dialog will appear. Set the parameters for the text you want to add, type your words into the field at the bottom of the dialog and click on **OK** to add the text. Please note that the fonts displayed by the **GIMP** are the fonts that are accessible to X on your system.

When you add text to an image, it is selected. If you move your mouse cursor over the text, you'll see the move icon, and you can click and drag to move the text to an appropriate spot on the image.



Pick colors from the image

The **color picker** is used to choose a color from the image. The **Color Picker** dialog will display the individual RGB values and the hexadecimal for the selected pixel. Additionally, the active color box at the bottom of the Toolbox will change to the value of the chosen pixel.



Fill with a color or pattern

The **bucket** or **bucket fill tool** is used to fill a selection with the active color or with a pattern.



Fill with a color gradient

The **blend** or **gradient tool** is used to fill a selection with a range of colors.

To add a gradient to a selection, first double-click on the gradient button to display the gradient **Tool Options**. The **Blend** pull-down menu controls the range of colors you're applying. You can blend from the foreground color to the background color or vice versa, you can blend from a color to transparency, or you can choose from the many custom gradients available in the **GIMP**.

You can set the "shape" of the gradient with the **Gradient** pull-down menu on the gradient **Tool Options**. You can also set whether the gradient repeats or not, using the **Repeat** pull-down menu.

Once you've chosen a gradient to use, click and drag on your selection, or on the entire image, to apply your gradient. For most gradients (but not all), the length of the drag affects the range of the gradient; a short drag will equate to a compact gradient with more defined edges between colors, while a long drag will produce a wider gradient with less clearly defined edges.



Draw sharp pencil strokes

The **pencil** is used to draw lines, curves or shapes with sharp, defined edges. See *The Brush Selection Dialog* in Section 9.1.1 for instructions on selecting brushes for use with the pencil.



Paint fuzzy brush strokes

The **paintbrush** is used to draw lines, curves or shapes with soft edges. See *The Brush Selection Dialog* in Section 9.1.1 for instructions on selecting brushes for use with the paintbrush.



Erase to foreground or transparency

Oddly enough, the **eraser** is used for erasing. See *The Brush Selection Dialog* in Section 9.1.1 for brushes that can be used with the eraser.



Airbrush with variable pressure

The **airbrush** is used to draw in various levels of opacity up to full (which would be like drawing with the paintbrush). See *The Brush Selection Dialog* in Section 9.1.1 for instructions on selecting brushes.



Paint using patterns or image regions

The **clone** or **stamp tool** is used to select portions of the image, which are then used as "paint." The clone tool can be used to remove unwanted portions of an image, by painting over them with patterns from the image itself.

Click on the clone tool in the Toolbox to activate it. Then hold down the [Ctrl] key and click on the place in the image which will supply the pattern for painting. Release the [Ctrl] key. Click and drag to apply the selected pattern to the image. A + sign will follow along with the pencil icon of the clone tool. The + sign starts at the spot you picked and indicates where the "paint" is coming from in the image.



Blur or sharpen

The **blur/sharpen tool** is used to paint a blurring or sharpening effect on specific portions of an image.



Draw in ink

The **ink pen** is used to draw lines similar to those drawn using a pen; the line's thickness depends upon the speed of the drawing device (a pen tablet or your mouse). The pen's **Tool Options** include settings for the size and shape of the pen.



Dodge or Burn

The **dodge/burn tool** is used to adjust intensity values (the amount of white in each pixel) in an image.

The **Tool Options** for the dodge/burn tool allow you to set whether you want to dodge (increase the intensity) or burn (decrease the intensity). The **Exposure** slider sets how much you want to change the intensity, from 0 to 100%.

The **Mode** sets which pixels will be affected: **Highlights**, pixels with high intensity values; **Midtones**, pixels with medium intensity values; or **Shadows**, pixels with low intensity values. The dodge/burn tool is used with a particular brush; see *The Brush Selection Dialog* in Section 9.1.1 for more information about brushes.



Smudge

The **smudge tool** is used to blur or smear an image with the currently selected brush. See *The Brush Selection Dialog* in Section 9.1.1 for more information on brushes.



Measure distances and angles

The **measure tool** can be used to measure a distance on an image, or to measure an angle from the horizontal.


Colors

At the bottom of the Toolbox, you'll see the active colors display:


Figure 9–2 The Active Colors



The active color square is depressed; you can see in Figure 9–2, *The Active Colors* that the black foreground color square is the active color.

By default, black is the active foreground color and white is the active background color. To "flip" the foreground and background colors, click on the .

Click on the active color box to display the **Color Selection** dialog:

Figure 9–3 The Color Selection Dialog

You can select a color "freehand," by choosing one of the radio buttons for hue, saturation, value, red, green or blue, picking a range of color from the middle bar, and then picking a specific color from the color square.

You can also just type in the correct RGB values or the hexadecimal value for the color you want to select.

As you choose colors by clicking on the color square, or by typing in exact values, you'll see the active color box on the Toolbox change to the new color.

Color selection is a lot more complicated than what is covered here. If you're interested, you should check out some of the more in-depth resources in Section 9.3, *Where to Find More Information*.

The Active Brush, Pattern and Gradient

A square on the bottom right of the Toolbox displays the active brush, pattern and gradient:

Figure 9–4 The Active Brush, Pattern and Gradient



If you want to choose a different brush, pattern or gradient, double-click on either the active brush, pattern or gradient to display the palette of choices for each. Click on a choice from the palette to make it the active brush, pattern or gradient.

Selecting Areas in an Image

The GIMP provides six selection tools. You can use these selection tools by clicking and dragging on the portion of the image you want to select. The GIMP also provides ways to modify your selection.


These modifications will only work while your selection is active. You can recognize an active selection by the moving dotted line around its perimeter (marching ants).

- When the mouse cursor is on your selection, the cursor becomes the move symbol. If you click and drag on the selection, you can move the selection.
- The [Shift] key allows you to add to the selection. Once you have an active selection, if you hold down the [Shift] key and then use a selection tool to select a portion of the image, the new selection will be added to the active selection.
- The [Ctrl] key allows you to subtract from an active selection. If you hold down the [Ctrl] key and then use a selection tool to select a portion of an active selection, that portion will be removed.

The Brush Selection Dialog

To display the **Brush Selection** dialog, click on the active brush on the Toolbox. The **Brush Selection** dialog looks like Figure 9–5, *The Brush Selection Dialog*:

Figure 9–5 The Brush Selection Dialog



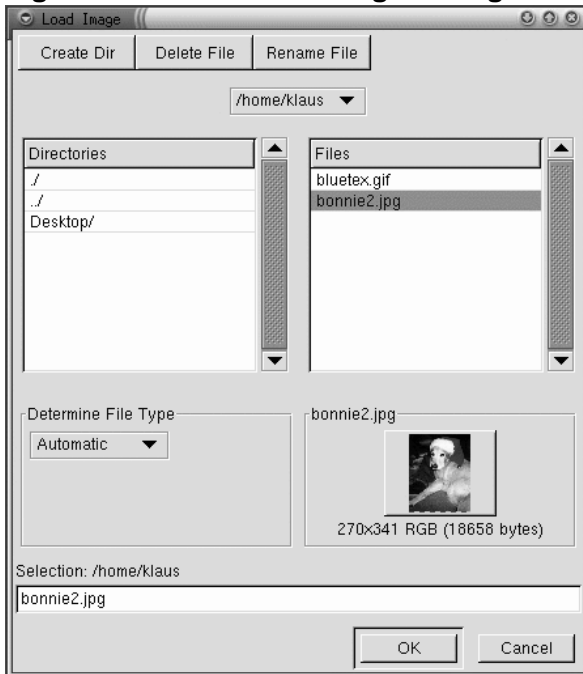
When you use GIMP drawing tools, you can modify the attributes of the line you're drawing by selecting a particular brush.

Click on a brush in the palette to select it. If the brush is larger than the square can display, click and hold to see the entire brush.

9.1.2 Loading a File

Before you can do any image manipulation, you need to somehow acquire an image. For example, you can scan an image, you can find one on the Internet (be careful about copyright restrictions!), or you can take a picture with a digital camera.

To load an existing file, click on **File => Open**. You'll see the **Load Image** dialog, as shown in Figure 9–6, *The Load Image Dialog*.

Figure 9–6 The Load Image Dialog

The **Load Image** dialog displays your working directory when you started GIMP. You can navigate up and down the filesystem tree by double-clicking on the **Directories** list on the left, then select a file to open from the **Files** list on the right.

Filename completion is supported by the GIMP. If you type the first letter (or more) of a filename into the **Selection** field and press the [Tab] key, the view will change to only those subdirectories and/or files beginning with that letter or letters.

The file you select will appear in the **Selection** field near the bottom of the dialog. A thumbnail preview will be displayed on the dialog; alternatively, you'll see a **Generate Preview** button. If you want to see a thumbnail of the image, click on the **Generate Preview** button.

Once you've selected a file, click on the **OK** button to open it. You can also double-click on a filename to open the file.

9.1.3 Saving a File

To save an image file, right click on the image and choose **File => Save** (or **Save as**). You'll see the **Save Image** dialog if you choose **Save as** or if you choose **Save** but the file has not been saved before.

The **Save Image** dialog looks almost exactly like the **Load Image** dialog and navigation of the filesystem tree and choosing files works in the same way.

When you're saving an image, you'll need to choose what format to save it in. The GIMP supports a wide variety of image formats. In this chapter, we're going to talk about three of them: `.gif`, `.jpg` and `.xcf`.

Image Formats


The vast majority of images on the Web are in one of two formats: `.gif` or `.jpg`. Both formats utilize compression to provide smaller file sizes and shorter download times (download speed is of paramount importance for everything on the Web). In most other ways, however, these two image formats are very different, and they're best used for different purposes.

The `.jpg` format is best used for photographs or images which use a lot of different colors.

The `.gif` format is best used for images which only use a few colors or which incorporate relatively large areas of the same color.

Before the GIMP will allow you to save an image as a `.gif`, you'll need to convert to indexed color or to grayscale. Convert an image to indexed color by right-clicking on the image and then selecting **Image => Mode => Indexed**. The **Indexed Color Conversion** dialog will be displayed:

Figure 9–7 The Indexed Color Conversion Dialog



To minimize file size, a `.gif` image can only display a maximum of 256 colors. When you convert to **indexed color**, the RGB colors in the image are "standardized" to a particular palette of colors. If your image is going to be used on the Web, the safest palette choice is **Use Custom Palette: Web**. The palette optimized for the Web includes the group of 216 colors that are used by the majority of Web browsers.

After you've converted your image to indexed format, the **GIMP** will allow you to save it as a `.gif` file.


If you're working on an image in the **GIMP**, and you want to save it in order to keep working on it later, you should save it as a `.xcf` file. The `.xcf` format is the **GIMP**'s native file format. If you create an image in the **GIMP**, or if you open an image and add **GIMP** effects (such as layers or masks) to it, you have added information to the image file. This information will be lost if you save the image as a `.gif` or `.jpg`.

If you save a file as a `.xcf` file, you'll preserve that information. When you re-open the `.xcf` file, you'll be able to edit the file as if you had never closed it.

9.1.4 The Image Window

When you load an image or open a new image, the GIMP surrounds the image with a frame which includes helpful information and tools, as shown in Figure 9–8, *The Image Window*.

Figure 9–8 The Image Window



At the top, the image window displays a file identifier consisting of:

- the name of the file (preceded by an * if the file has been modified)
- the window number (sequential from 0 for the first image opened when you start the GIMP)

- the view number (sequential from 0 for each image view)
- the image type -- RGB, Indexed, or Grayscale
- the amount of the image that can be seen in the window, depending on if you've zoomed in on the image

At the top left, if you click on the  icon, the image menu will appear.

To help you visualize the size of your image, the **GIMP** provides horizontal and vertical rulers. You can toggle these off and on by right-clicking on the image and selecting **View => Toggle Rulers**.

At the bottom left, you'll see the position of the mouse cursor on the image, measured from **0,0** at the top left corner.

Just above the mouse cursor position window, you'll see two square icons, which are used for creating quick masks. Basically, masks are selections which can be saved and used again. We won't cover masks in this elementary chapter, but if you use the **GIMP** for complex images, you'll find masks to be very helpful. Please see some of the sources listed in Section 9.3, *Where to Find More Information* for more information on masks.

Near the bottom right, you'll see an icon that looks like the move tool icon. This is called the **panner**. If you're zoomed in on an image, you can click on the panner to display a thumbnail of the whole image. Then you can click and drag on the thumbnail to pan the image view in your image window.

Along the bottom of the window, next to the mouse cursor position window, you'll see two status fields and then a **Cancel** button. The first status field displays the file name, unless a plug-in is running (for example, if you're applying a filter or a Script-Fu to an image). If a plug-in is running, the name of the plug-in will appear in the first status window. The next field is empty, but will display a status bar for a running plug-in. Finally, the **Cancel** button allows you to stop a running plug-in.

9.1.5 Helpful Keyboard Shortcuts

Like many other applications, the GIMP makes it easy to accomplish tasks without taking your hands off the keyboard. Here are a selection of helpful keyboard shortcut commands:

Table 9–1 Keyboard Shortcut Commands

Command	Keystroke	Command	Keystroke
copy	[Ctrl]-[c]	select all	[Ctrl]-[a]
cut	[Ctrl]-[x]	select none	[Ctrl]-[Shift]-[h]
invert selection	[Ctrl]-[i]	undo	[Ctrl]-[z]
paste	[Ctrl]-[v]	zoom in	[=]
save file	[Ctrl]-[s]	zoom out	[-]

9.2 Using the GIMP to Create a Web Page


The easiest way to learn about anything is to try it yourself. In this section you'll learn the steps needed to create some simple Web page graphics with the GIMP. The GIMP can do very complicated image manipulations -- enough for a hefty book. In this section, we really will be just scratching the surface of its abilities.

9.2.1 Enhancing a Picture

For this short example, we'll take a single picture and:

- perform automatic color correction
- correct a common problem in photographs -- flash-reflecting eyes
- remove some unwanted portions of the picture


First, a scanned image is loaded into the GIMP, using **File => Open**:

Figure 9–9 The Original Scanned Photograph, Including Flaws

As you can see, the original photograph of Bonnie the dog isn't very good. The overall range of colors seems faded. Her eyes are an obvious problem -- they're reflecting the flash in a scary green color. Also, two distracting areas could be removed: the dark smudge at the top of her leg, and the small amount of a person's hand that can be seen by Bonnie's face.

First, the **GIMP** will be used to do some automatic color correction. Right-click on the image and select **Image => Colors => Levels** to display the **Levels** dialog. Since we're interested in color improvement, select **Red** from the **Modify Levels for Channel** pull-down menu, as seen in Figure 9–10, *The Levels Dialog*.

Figure 9–10 The Levels Dialog



The **histogram** displays the red values for all of the pixels in the image. The range of red values in the picture doesn't spread over the entire available tonal range (the lower bar, which shades from black to red). In other words, none of the pixels in the image have red values in the darkest range -- the histogram is flat in the very dark region of the range.

If you look at the **Blue**, **Green** and **Value** (brightness) levels, their histograms look similar. Overall, the darkest colors in the image are not as dark as they could be, given the range of colors. The image will look better if the range of values is compressed, so that the range of values in the image spreads over the entire range of values that are available.


This sounds complicated, but the **GIMP** makes it easy. Just click on the **Auto** button in the **Levels** dialog. **Auto** will automatically adjust all of the color ranges (red, blue

and green). If you would like to also adjust the **Value** (brightness) of the image, you'll need to select **Value** from the **Modify Levels for Channel** pull-down menu and click on **Auto** again.

If the **Preview** checkbox on the **Levels** dialog is checked, then you'll see a preview of the correction on the image. If you're happy with it, press **OK** on the **Levels** dialog to accept the changes.

Now, the Bonnie image looks like Figure 9–11, *Image After Automatic Color Levels Adjustment*.

Figure 9–11 Image After Automatic Color Levels Adjustment



Next, we'll use the clone tool to remove distracting parts of the image. The clone tool grabs a portion of the image, and then allows you to paint the image with that portion. This allows you to replace unwanted artifacts in the image with a realistic-looking pattern -- a pattern actually found in the image.

First, you can use the [=] key to zoom in on the unwanted dark smudge on Bonnie's leg. Then select an appropriate brush. (To select a brush, click on the active brush in the Toolbox and select one from the **Brush Selection** dialog.) For this purpose, the **Circle (07) (7 x 7)** was chosen.


Select the clone tool from the Toolbox. With the [Ctrl] key pressed, click on an area to the left of the dark smudge, to select an area to apply as a pattern. Release the [Ctrl] key. The crosshairs icon changes into a pencil icon, and you can paint using the selected area as a pattern.

As you paint using the clone tool, you'll see the crosshairs icon follow along at a distance with your paintbrush. The crosshairs icon marks the area which is providing the pattern for the paintbrush.

The same technique can be used to remove the small portion of a person's hand seen by Bonnie's face. Use the same brush and select an area from the blue jeans near the hand to clone.

The **GIMP** can be used to fix the eerie green reflection in her eyes. From the **Brush Selection** dialog, select the **Circle Fuzzy (11) x 11)** brush. Use the color picker tool to select a black color from the non-green edge of her eye. Then use the paintbrush tool to paint in the green area -- the most difficult part is leaving a tiny spot of white in each eye, to make the eyes look a little more natural (not completely flat black).

After automatic color levels adjustment and removal of the smudge, the hand, and the green reflection, the image looks like Figure 9-12, *Image After Removal of Unwanted Portions*.

Figure 9–12 Image After Removal of Unwanted Portions

Finally, the edge of the image will be blurred, purely as a decorative effect.

Create a new image, larger than Bonnie's photograph, with a white background. Cut and paste the Bonnie image into the larger, white image. Use the rectangular select tool to select a rectangle just slightly smaller than Bonnie's image. Then right-click on the image, and choose **Select => Invert** to invert the selection, so a rectangular "frame" around the image, is selected, as in Figure 9–13, *A "Frame" Selection Around the Image*.

Figure 9–13 A "Frame" Selection Around the Image



To blur the edge, right-click and select **Filters** => **Blur** => **Gaussian Blur (IIR)**. With both of the **Blur Radius** settings at 10, click on the **OK** button to apply the filter.

Then the image is cropped to leave just a small white border, so it now looks like Figure 9–14, *The Final Image*.

Figure 9–14 The Final Image

Finally, the image is saved as a .jpg file for use on our Web page.

9.2.2 Creating a Web Page Background Graphic


Since the Web page displays a picture of a dog, a stylized pawprint will be used as the background.

The first step is to make a black and white pawprint image. Open a new image window using **File => New**. In the **New Image** dialog that appears, set the **Width** and **Height** both to 500 pixels, to give you some room to work -- you can always crop the image once you're finished. Set the **Fill Type** to **White** and click on the **OK** button.

To make the image, create an oval selection, then use the bucket tool to fill in the selection with black. Copy the oval selection ([Ctrl]-[C]) and paste it ([Ctrl]-[V]) twice as you move the three ovals into an appropriate pattern. Then, use the paintbrush tool

with a small brush to draw the pad of her foot. When you're drawing something, it is easier to draw a large image, then scale it after you're finished. The original black and white image, sized 128 x 134 pixels, looks like Figure 9–15, *A Pawprint in Black and White*.

Figure 9–15 A Pawprint in Black and White



The pawprint image is the basic building block for the next image -- a series of tracks that "walk" across a page, as seen in Figure 9–16, *Tracks in Black and White*. To create the tracks image, the pawprint image is:

- Scaled -- right-click on the image, then click **Image** => **Scale Image** to display the **Scale Image** dialog. The image was scaled by **Ratio** .85, because it was a little too large for the desired effect.
- Rotated -- choose the transforms tool from the Toolbox and rotate the image freehand to about a 50-degree angle from horizontal
- Copied and pasted into a new larger image -- using [Ctrl]-[C] and then using [Ctrl]-[V].

The first pawprint is pasted into the new image with a [Ctrl]-[V]. The pawprint is then moved into the desired position. After it is placed correctly, another copy of the pawprint is pasted into the image. The second pawprint is now the active


image, pasted directly over the first image. The second pawprint is then moved off the first pawprint into an appropriate spot. The same steps were repeated for the third and fourth pawprint. If you move something into the wrong spot, but you've already selected something else, you can always use [Ctrl]-[Z] to back up and try again.

Next, a **filter** will be used to create a special effect. Filters provide a particular look or feel to an image. Right-click on the image and look at **Filters** to see a list of **GIMP** filters. For this image, we'll first use the **Gaussian Blur (IIR)** filter and then the **Bump Map** filter.

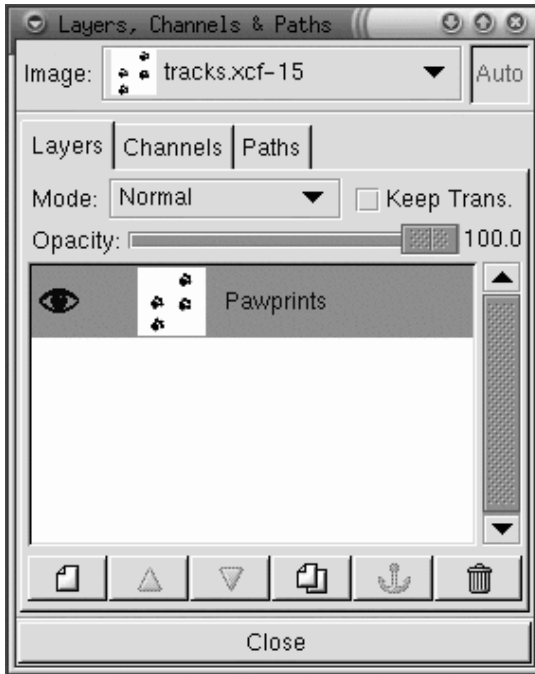
After the pawprints are placed in the right spots, the image is blurred with **Filters => Blur => Gaussian Blur (IIR)** set to both a horizontal **Blur Radius** of 5, the defaults. The image is blurred because the next filter, **Bump Map**, acts on light and dark pixels, so its results will appear more natural with a bit of shading between the black and white of the original image.

The black and white tracks image, blurred slightly and scaled again to 50% so that it measures 250 x 250 pixels, looks like Figure 9–16, *Tracks in Black and White*.

Figure 9–16 Tracks in Black and White




Next the **Cork board** pattern is added to a new **layer** of the image. Use **Layers => Layers, Channels & Paths** to display the **Layers, Channels & Paths** dialog, as shown in Figure 9–17, *The Layers Dialog*.

Figure 9–17 The Layers Dialog

Layers keep portions of your working image separate from each other, so you can work on one piece of the image without affecting the other. Layers are essential for manipulating complex images.

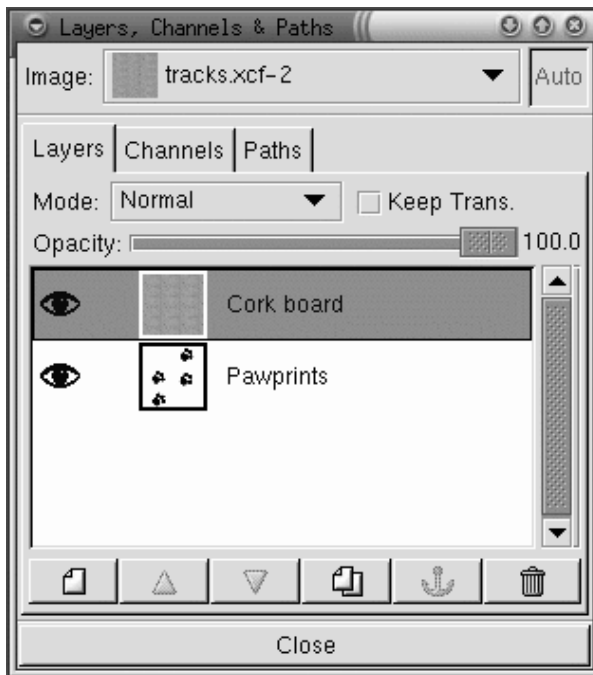
The **Layers** dialog displays all of the layers in the image. As shown in Figure 9–17, *The Layers Dialog*, the image only includes one layer, which was re-named to **Pawprints** (right-click on the layer, then choose **Edit layer attributes** and fill in a descriptive name).


Next, a new layer is added for the pattern by clicking on the **New layer** icon . On the **New Layer Options** dialog that appears, leave the **Layer Width** and **Height** at the default values (the layer will be the same size as the image). Set the **Layer Fill Type** to **White** and click on the **OK** button.

New layers are automatically created on top of old layers, so you won't see the original layer in the image.

Next, the layer will be filled with the **Cork board** pattern. To fill a layer with a pattern, make sure that the layer is "active" (highlighted with a colored bar on the layers dialog -- inactive layers will be white). See Figure 9–18, *The New Layer in the Layers Dialog* for an example; the **Cork board** layer is active. Then, on the **Tool Options** for the bucket tool (double-click on the bucket tool to display the **Tool Options**), select **Pattern Fill**. Click on the active pattern on the Toolbox to see the patterns palette and select **Cork board**. Then use the bucket tool to fill the **Cork board** layer.

Figure 9–18 The New Layer in the Layers Dialog




On the **Layers** dialog, make the pawprints layer the active layer by clicking on it. Then click on the eye icon () to make the cork board layer invisible. (Another click on the eye icon's spot will make the layer visible again.)

On the image, which should just show the pawprints image as seen in Figure 9–16, *Tracks in Black and White*, use the magic wand tool to select a white pixel. Then right-click on the image and go to **Select => Invert** to invert the selection, so only the black or gray pixels are selected.

Now when you click on the eye icon on the layers dialog, to display the cork board layer, you'll see the pawprint selection, as seen in Figure 9–19, *The Pawprint Selection*.


Figure 9–19 The Pawprint Selection



The next step is to darken just the selected pawprints area. The pawprints need to be darkened so that the **Bump Map** filter will affect the pawprints, as well as the rest of the cork board pattern. The **GIMP** provides a number of different ways to darken the selected area. One simple way is to right-click on the image (with the selection still active), then select **Image => Colors => Brightness-Contrast**. On the **Brightness-Contrast**


dialog, the **Brightness** slider is moved to **-40**, as shown in Figure 9–20, *The Brightness-Contrast Dialog*.

Figure 9–20 The Brightness-Contrast Dialog



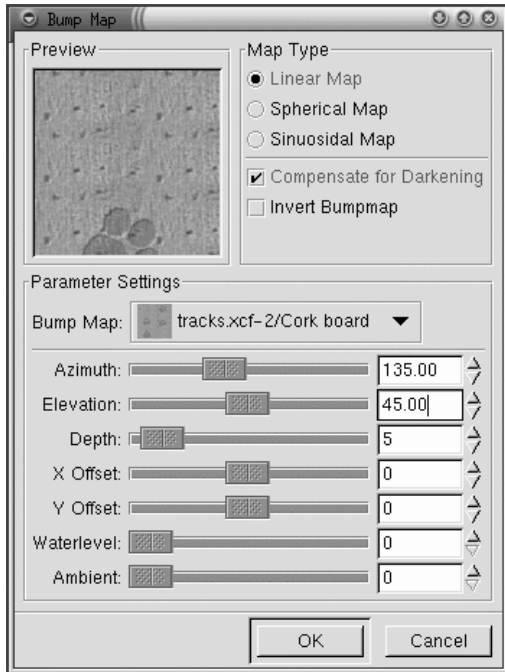
Click on the **OK** button to darken the selected areas. Then right-click on the image, and choose **Select => None** (or use the keyboard shortcut [Shift]-[Ctrl]-[A]) to remove the selection, so you'll see Figure 9–21, *The Pawprints on the Cork Board Background*.

Figure 9–21 The Pawprints on the Cork Board Background




To apply the filter, right-click on the image, then select **Filters** => **Map** => **Bump map** to display the **Bump Map** dialog:

Figure 9–22 The Bump Map Dialog



On the **Bump Map** dialog, the **Compensate for Darkening** checkbox is checked and the **Depth** is increased to 5 (which increases the perceived depth of the pawprint). A preview of the image after the filter is applied is provided to show how the different **Bump Map** parameters can affect the appearance of an image. After selecting **OK**, the **Bump Map** filter is applied to the image as seen in Figure 9–23, *The Final Background Image*.

Figure 9–23 The Final Background Image

Background graphics for Web pages need to **tile** -- any pattern on the graphic needs to match up on all sides so that when they're repeated on the page, the background appears to be seamless. In this case, since the cork board pattern is already tiled, there is no need to re-tile it. However, if you're working on a non-tiled graphic, the GIMP will tile a background for you using the **Filters => Map => Tile** filter.

9.2.3 Creating Title Images for Web Pages

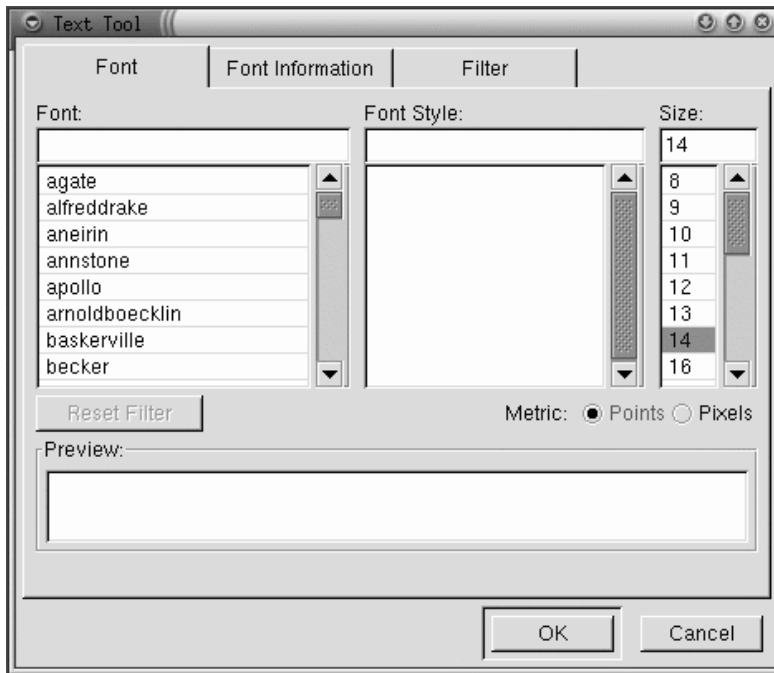
The GIMP is very helpful for creating logos (text for titles and links), buttons, and bullet graphics for Web pages.

In this example, we'll use a couple of GIMP filters to add a custom look to some text for our Web page example.

The first step is to open a new file, sized 400 x 100 pixels, with a white background. Use **File => New**, add the correct **Width** and **Height** parameters, and select the **Transparent** radio button.

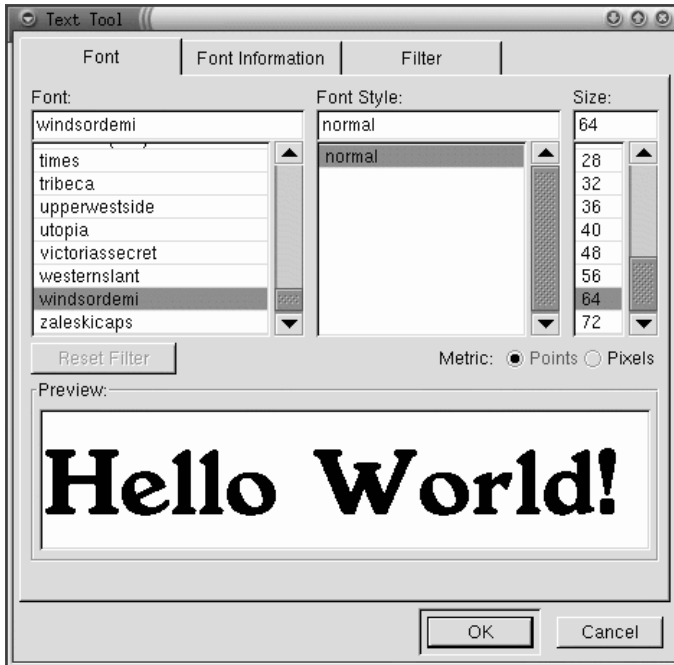
Click on the text tool in the Toolbox to select it. Then click on the image, and the **Text Tool** dialog will appear as in Figure 9–24, *The Text Tool Dialog*.

Figure 9–24 The Text Tool Dialog




Type the text you want into the **Preview** field near the bottom of the dialog. Then select a **Font**, **Font Style**, and **Size** from the pull-down menus in the dialog. The text you've typed in will change to preview your selections, as shown in Figure 9–25, *Previewing Text*.

Figure 9–25 Previewing Text



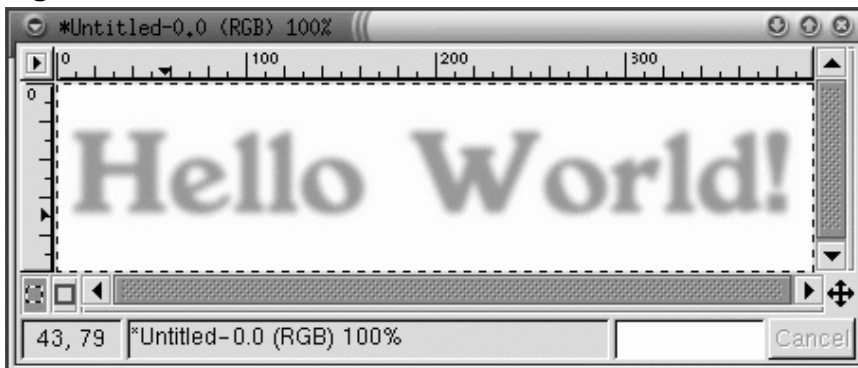
This text is actually going to be the drop shadow for the "real" text. Change the foreground color to a color for the drop shadow. For this example, a medium gray, with an RGB value of 153 153 153, is selected. Once you click on **OK** in the **Text Tool** dialog, the text will appear in the working window, as shown in Figure 9–26, *Text in an Image Window*.

Figure 9–26 Text in an Image Window



Click on the rectangle selection tool, then click on the image to place the text selection onto the image. To create a drop shadow, right-click on the image, and apply **Filters => Blur => Gaussian Blur (IIR)**. The **Gaussian Blur** dialog allows you to set the **Blur Radius**, which was left at the default of 5 pixels to produce the blurred effect in Figure 9–27, *Blurred Text*.


Figure 9–27 Blurred Text



Now the "real" text should be added to the image. First, the active foreground color needs to be changed. In this case, a faded red (RGB 204 051 051) is chosen. Then with the text tool active, click on the image. The **Text Tool** dialog will appear. Make sure the same settings are selected (they will be there by default, unless you closed

the GIMP). Click on **OK** to add a red "Hello World!" to the image, and then move it into the right position, as shown in Figure 9–28, *Colored Text with a Drop Shadow*.

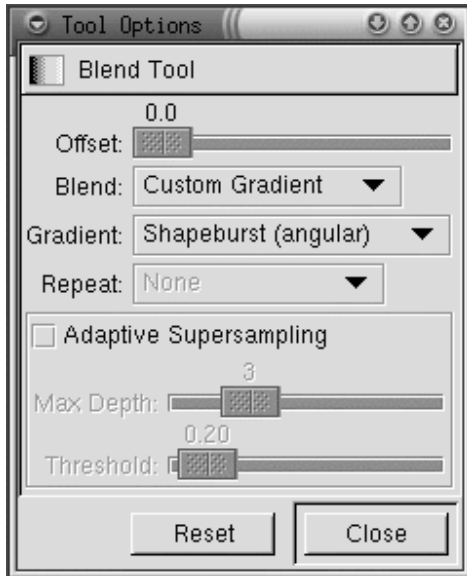
Figure 9–28 Colored Text with a Drop Shadow



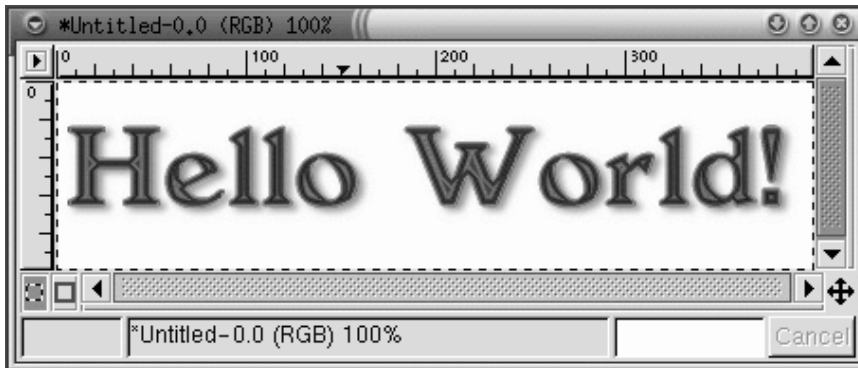
Perhaps this still looks a little plain, and you'd like something a little more showy for your Web page. One option is to apply one of the GIMP's gradients to the text.

With the red "Hello World!" as the active selection, double-click on the gradient tool to make the gradient tool active and display its **Tool Options**. For this effect, **Custom Gradient** is selected as the **Blend** and **Shapeburst (angular)** is selected as the **Gradient**, as shown in Figure 9–29, *Gradient Options*.

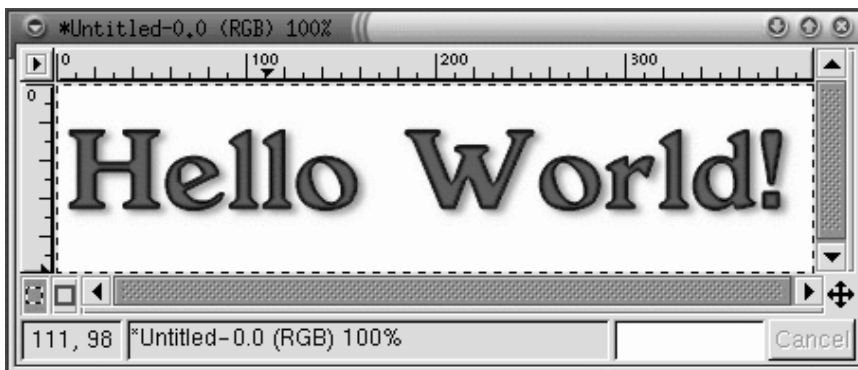
Figure 9–29 Gradient Options



The next step is to choose which of the GIMP's many gradients to use. Click on the active gradient on the Toolbox to display the palette of available gradients. Click on a gradient to make it the active gradient. To apply a gradient to your text, make sure the text is still the active selection. Then click and drag on part of the text (since we chose **Shapeburst (angular)**, the length of the drag won't affect the appearance of the gradient, which will instead follow the edges of the selection). For example, if you apply these settings with the **Greens** gradient, the image looks like Figure 9–30, *Text in the Greens Gradient*.

Figure 9–30 Text in the Greens Gradient

Starting with the plain red "Hello World!" (just use [Ctrl]-[Z] to go back) and the same gradient parameters, the **Tube_Red** gradient produces this effect, which will be used in the final Web page, as shown in Figure 9–31, *Text in the Tube_Red Gradient*.

Figure 9–31 Text in the Tube_Red Gradient

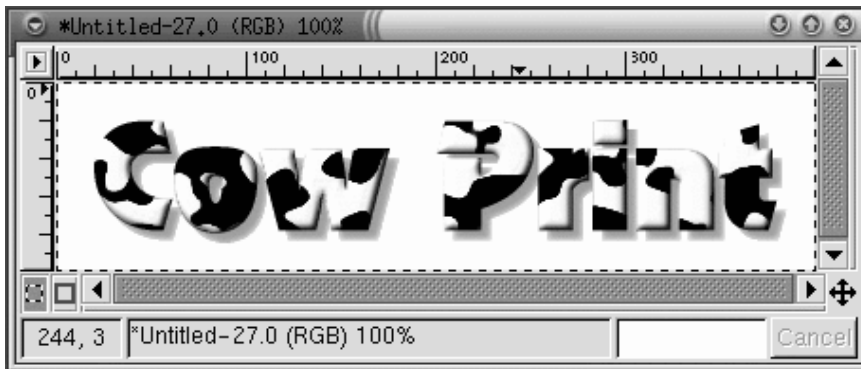
Script-Fu

As we've shown, you can use the tools and filters of the GIMP to create unique logos. You can also use the GIMP's **Script-Fu** to make graphics that may not be completely unique, but that are quickly created, easy, and very professional-looking.

Script-Fus provide a way for non-programmers to take advantage of GIMP's scripting abilities. Basically, they provide different GUIs for creating logos, patterns, brushes, and other special effects. You can also create your own Script-Fus. See Section 9.3, *Where to Find More Information* for sources of information about writing your own Script-Fus.

The easiest way to get an idea of what Script-Fu is all about is to try one. From the Toolbox, click on **Xtns => Script-Fu** to see the menu of available Script-Fus. For example, if you chose **Xtns => Script-Fu => Logos => Bovination**, you'll see the **Script-Fu: Logos/Bovination** dialog window. You can change the **Script Arguments**; in this case, all we changed was the **Text**. Once you click on **OK**, you'll see the GIMP apply a series of filters, and then a new image window will display the graphic created by Script-Fu: Figure 9-32, *Cow Print Text*.

Figure 9-32 Cow Print Text



A cow print logo may not be right for your Web page, but it does provide an excellent example of Script-Fu's ability to create special effects!

9.2.4 Combining the Graphics We've Created on a Web Page

The background, title and edited photograph can all be used on the same simple Web page: Figure 9-33, *Sample Web Page*.

Figure 9–33 Sample Web Page

9.3 Where to Find More Information

The GIMP can do so much more than what we've covered in this chapter! Please seek out some of the following resources if you're interested in learning more about the GIMP.

Many books have been written about the GIMP. If you need in-depth information about the many capabilities of the GIMP, try your favorite bookstore. The following books were available at the time of this writing:

- *The Artists' Guide to the GIMP* by Michael J. Hammel
- *GIMP Essential Reference* by Alex Harford
- *GIMP for Linux Bible* by Stephanie Cottrell Bryant, et al

- *GIMP: The Official Handbook* by Karin Kylander and Olof S. Kylander (including "Mike Terry's Black Belt School of Script-Fu," an excellent how-to chapter on writing Script-Fus)
- *Grokking the GIMP* by Carey Bunks
- *Teach Yourself GIMP in 24 Hours* by Joshua and Ramona Pruitt

Information about the GIMP is also available on the Web:

<http://www.gimp.org/>

The official GIMP website

<http://www.rru.com/~meo/gimp/faq-user.html>

A Frequently Asked Questions (FAQ) list for questions commonly asked about the GIMP by GIMP users (as opposed to developers)

<http://manual.gimp.org/manual/>

The *GIMP User Manual* (by Karin Kylander and Olof S. Kylander) website, which includes the chapter "Mike Terry's Black Belt School of Script-Fu," an excellent resource for anyone who wants to learn how to write a Script-Fu

<http://empyrean.lib.ndsu.nodak.edu/~nem/gimp/tuts/>

The GIMP Tutorials Pointer Page

<http://gimp-savvy.com/>

The companion website to the book *Grokking the GIMP*, by Carey Bunks

<http://tigert.gimp.org/gimp/>

The GIMP website of tigert (Tuomas Kuosmanen), the artist who created Wilber (the GIMP mascot), as well as many other GIMP graphics

<http://brahms.fmi.uni-passau.de/~anderss/GIMP/>

The GIMP 1.0 Quick Reference Sheet -- a tri-fold list of keyboard shortcuts, etc.

Part II Managing Files and Directories

10 File Managers in GNOME and KDE

In both GNOME and KDE, the file managers are true workhorses. Use these powerful applications to help you create, edit, and delete files and directories, as well as accomplish other tasks throughout your system.

Both file managers support drag and drop, and provide you with a fast and easy way to "see" your system. Either file manager will open in your login directory by default. To view other directories on your system, click on the "file folder" representing the directory in the **tree view** pane on the left of the file manager. The contents of the folder will appear in the **directory window**, the pane on the right.

Do you have permission?

Unless you're root, you won't be able to gain access to all the files and directories on your system. If you don't have the permission to open, delete or execute a file, you'll receive an error message that your access is denied.

This chapter will provide an overview of GMC and KFM, the file managers for GNOME and KDE. To learn more details about using them, read the *GNOME User's Guide* or the KDE user's documentation in their respective help browsers. You'll also be able to find the latest documentation at their websites: <http://www.gnome.org> and <http://www.kde.org>.

10.1 The GNOME File Manager

Fast, friendly and flexible, the GNOME File Manager lets you easily browse and work with files and directories. The proper name for the file manager is **GNU Midnight Commander**, or **GMC**, because it's based on the Midnight Commander file manager.

You can still use Midnight Commander

If you want to use Midnight Commander, you can start it in an Xterm by typing `mc` at the prompt.

10.1.1 Using the File Manager

There are two ways to easily start GMC:

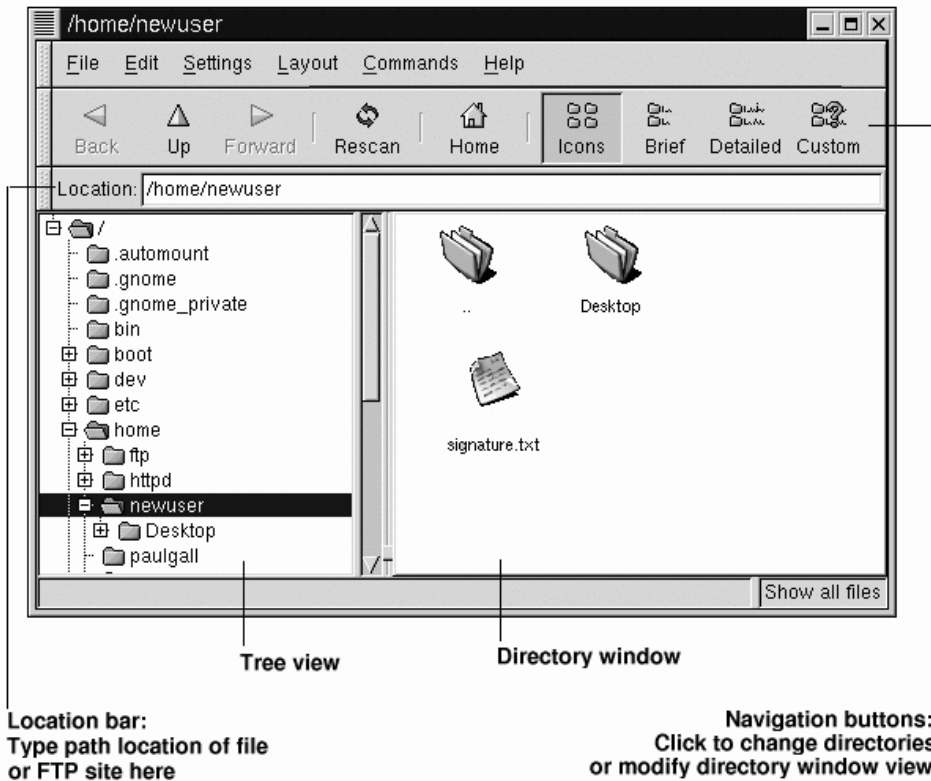
- From the **Main Menu Button** -- Click on **File Manager** entry.
- From the desktop -- Double-click on the file folder labeled **Home directory**, which



looks like:

When you first open the GMC, you'll see something similar to Figure 10-1, *A Sample GMC Window*.

Figure 10–1 A Sample GMC Window



As a new user, it's unlikely you'll find many new files or directories when you first open the file manager. That's because by default, GMC opens in your login directory. Note the file `signature.txt`, however, in Figure 10–1, *A Sample GMC Window*. This is the signature file the user `newuser` created while configuring Netscape Messenger.

The tree view, to the left, will reveal all of the directories on your system. Single-click on a folder to reveal its contents in the directory window, at right. Folders which have

a **+** next to them indicate that they hold subdirectories and the view is collapsed; click on the **+** to expand the view and reveal the subdirectories.

Use the following buttons on the menu above the **Location** bar to determine the detail of view you want:

- **Icons** -- The default view. Displays icons to represent each file and directory type.
- **Brief** -- Shows directories and files in a list, with directories at the top.
- **Detailed** -- Shows a list, along with details such as the size and modification time of the files.
- **Custom** -- Shows a customized detailed list (you can select items you want to include from the **Settings** menu, in **Preferences => Custom View**).

Change the order

If you view directory contents in one of the list choices, you can arrange your view of the list by clicking on one of the list headings. For example, if you want to see the contents alphabetically, click on the **Name** heading.

Navigation

Like many file managers, you can move around by using the scrollbar next to the tree view to locate a directory you want to move to.

You can also use the navigation buttons above the **Location** bar similar to the way you would use the navigation buttons in **Netscape Navigator**:

- **Back** and **Forward** -- Move you up or down through the directory listing history
 - **Up** -- Moves you "up" the directory tree
 - **Rescan** -- Refreshes the view of the current directory
 - **Home** -- Lets you jump back to your default directory
-

The **Location** bar allows you to type the directory location. If you know the path where you want to go, for example `/etc/X11`, you can type it into the **Location** bar and press [Enter] to jump to that location.

Actually, the **Location** bar can let you browse more than just your local system. If you're online, and you type in the location of an FTP site, such as `ftp.redhat.com`, you'll be able to see the files on that site which are available for download. The FTP command should look like this: `ftp://ftp.redhat.com`. To download, just copy a file to a directory on your local machine.

GMC plays well with Gnome-RPM

Speaking of downloaded files, you can use **GMC** to perform RPM-related work on new software packages, such as those you download or find on a CD. Just right-click on the package, and **GMC** offers you options to install, verify, query, and more (remember, you must be root to install or uninstall software). You can also drag and drop packages from the file manager into **Gnome-RPM** for installation (refer to Section 12.3, *Installing New Packages* for more on this).

To learn more about **Gnome-RPM** and RPM, turn to Chapter 12, *Using Gnome-RPM and RPM*.

Copying and Moving Files


You can copy and move files from the directory window to your desktop or to another folder (that is, directory). If you want to move a file to your desktop, just "pick it up" by clicking on it once and holding down the mouse button. Release the file on the desktop; by default, this action causes the file to be moved to your desktop.

If you'd prefer to copy the file rather than move it, click on the file, then drag it while holding down the [Shift] key.

Choose when you get there

Not sure whether you want to move or copy a file? Use the [Alt] key when you drag a file. A pop-up menu will appear once you release the file, letting you choose whether to copy or move the file. You can also choose to **link** the file, which essentially is a shortcut to the file in its original location.

Figure 10–2 Copying File Dialog



If you know the path to the directory to which you want to copy a file, you can right-click on the file and choose **Copy** from the menu of options which opens. A dialog will open, in which you can type the location where you want to copy the file (as in Figure 10–2, *Copying File Dialog*).

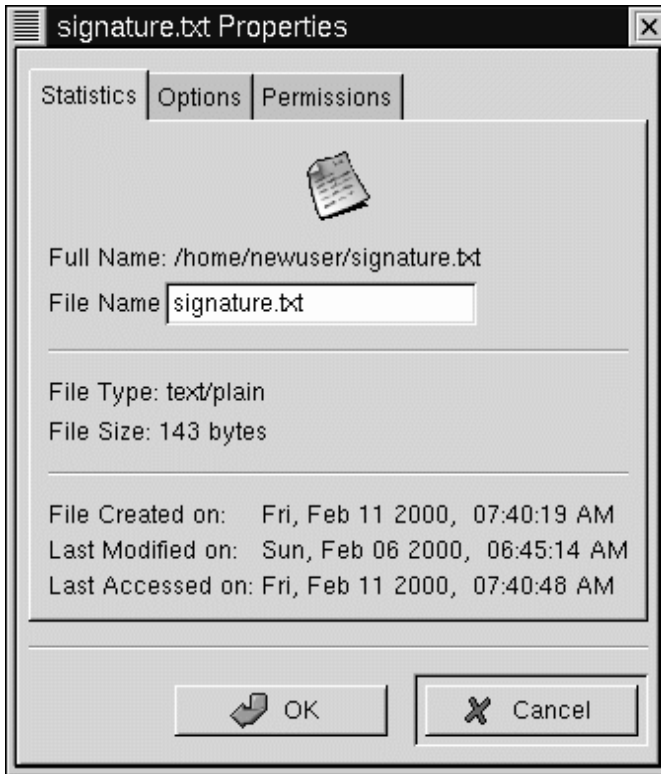
In the **Advanced Options** tab of the **Copy** dialog, you can choose to **Preserve symlinks**, which will make certain that if the file you're copying is a **symbolic link**, or shortcut, it will be the same shortcut after it's copied. Rather than risk losing a shortcut, you should make certain you select this option if you're copying -- or moving -- a file that is a shortcut to another file.

Moving works the same way

Just as you can copy a file by right-clicking, so too, you can move a file by right-clicking on it and selecting **Move** from the menu which appears. If you want to rename the file while moving or copying it, type its new name in the **Copy** or **Move** box which opens. For example, if you want to copy the file `signature.txt` to `/home/newuser/tigger` and rename it `backup.txt`, you would type `/home/newuser/tigger/backup.txt`.

File Properties


To change a file's properties, right-click on a file or directory in the directory window and choose **Properties**. A dialog, similar to Figure 10–3, *File Properties Dialog*, will open.

Figure 10–3 File Properties Dialog

In the **Statistics** tab, you'll find information about the file, such as the date and time it was created and modified.

In the **File Name** field, you can change the name of the file.

In the **Options** tab, you can change **action options** for the file, such as how the file will be opened, viewed and edited. You probably won't have a reason to worry about these items; they already have default settings.

Figure 10–4 The Permissions Tab in File Properties

In the **Permissions** tab (see Figure 10–4, *The Permissions Tab in File Properties*), you can change the permissions and ownership of a file (that is, if you have the permission yourself). You can change the read, write and execute settings, as well set **UID**, **GID** and **Sticky bits**.

UID, GID and Sticky bits?

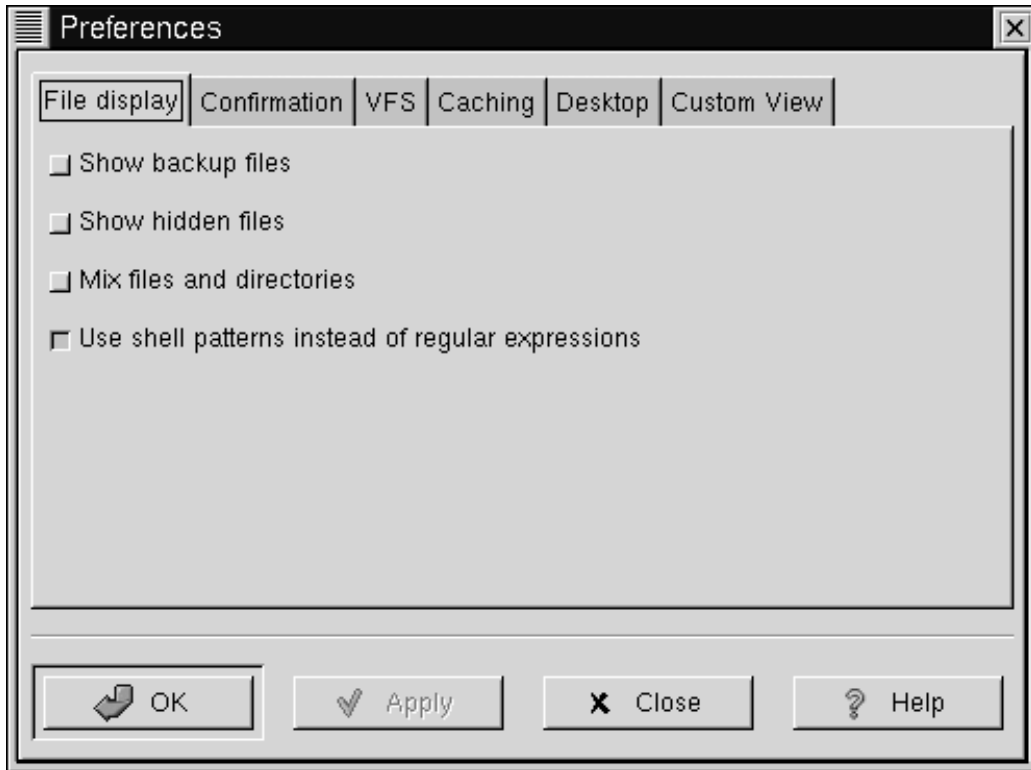
Essentially, these items refer to user management and permissions. UID means **User ID**, a unique number given to every user on the system; GID means **Group ID**, the ID of the user's default group. Sticky refers to **Sticky bit**. You probably won't have to worry about setting the Sticky bit.

For more about how to work with permissions, see Section 15.12, *Actions*. Once you understand the idea behind file permissions, you may find using the **Permissions** tab a convenient way to change a file or directory's permission settings.

10.1.2 Setting GMC's Preferences

You can adjust quite a few settings for GMC in the **Preferences** dialog (**Settings => Preferences**); see Figure 10–5, *The File display Tab in Preferences* for an example of the **Preferences** dialog.

Figure 10–5 The File display Tab in Preferences



The dialog includes the following tabbed sections:

- **File display** -- Modify the way files are viewed. Select **Show backup files** to reveal backup files; **Show hidden files** to show **dot files**, usually configuration files and directories; **Mix files and directories** allows you to see files and directories in an order you select in the directory window; you don't really have to concern yourself with the **Use shell patterns...** choice, because it concerns sorting with **regular expressions**, which is probably more useful to advanced users.

- **Confirmation** -- Select whether you should be required to confirm any decision to delete, overwrite and execute files. You can also choose whether you want to see a progress indicator while the actions are being executed.
- **VFS** -- VFS means **Virtual File System**. VFS means that you can work with files which aren't on your local machine. Information in this tab applies to actions related to FTP sites. **VFS Timeout** refers to the time (in seconds) you can be connected without activity, before you're disconnected. The entry in **Anonymous FTP password** refers to the identification you want to be transferred for **anonymous logins** at FTP sites. (Typically, this is your e-mail address.) Enable the option to **Always use FTP proxy** if you're required to use a proxy to connect to FTP sites.
- **Caching** -- These settings may help to increase the speed of GMC. Select **Fast directory reload** to store directory information in memory to speed the load time. Using this option may not always result in an updated display, however. **Compute totals before copying files** will cause GMC to calculate the number of files you're copying so it can provide information to you about the task. **FTP directory cache timeout** is the length of time (in seconds) to hold information about an FTP site in memory. **Allow customization of icons in icon view** lets you change the icon style in the directory window. Note that checking this may decrease performance.
- **Desktop** -- Here, you can specify some ways that you view the items on your desktop. Click inside the **Icon position** box to specify how the icons should be aligned on your desktop. Select **Automatic icon placement** to arrange icons on your desktop automatically, rather than arranging them manually. Click **Snap icons to grid** if you want icons on the desktop to be aligned. **Use shaped icons** and **Use shaped text** allow the icons and related text labels to be semi-transparent on your desktop, rather than "blocky" icons and labels.
- **Custom View** -- In this dialog, you can specify the additional information you want to see when you select the **Custom** button from GMC's menu.

10.2 The KDE File Manager

If you're accustomed to using Windows or a Macintosh, you're going to feel comfortable using KDE. In large part, that's due to the KDE File Manager, known as KFM.

Like GNOME's file manager, KFM is a graphical tool you can use to view and work with directories and files.

Useful online and offline

Because KFM uses HTML to display information, it serves as a decent Web browser, as well as a browser for your system. Just type a Web address (or, URL) into the **Location** bar when you're online, and you'll jump to your chosen website.

Similar to the GNOME file manager, you can also open FTP sites with the **Location** bar.

10.2.1 Using KFM

The easiest ways to start KFM are:

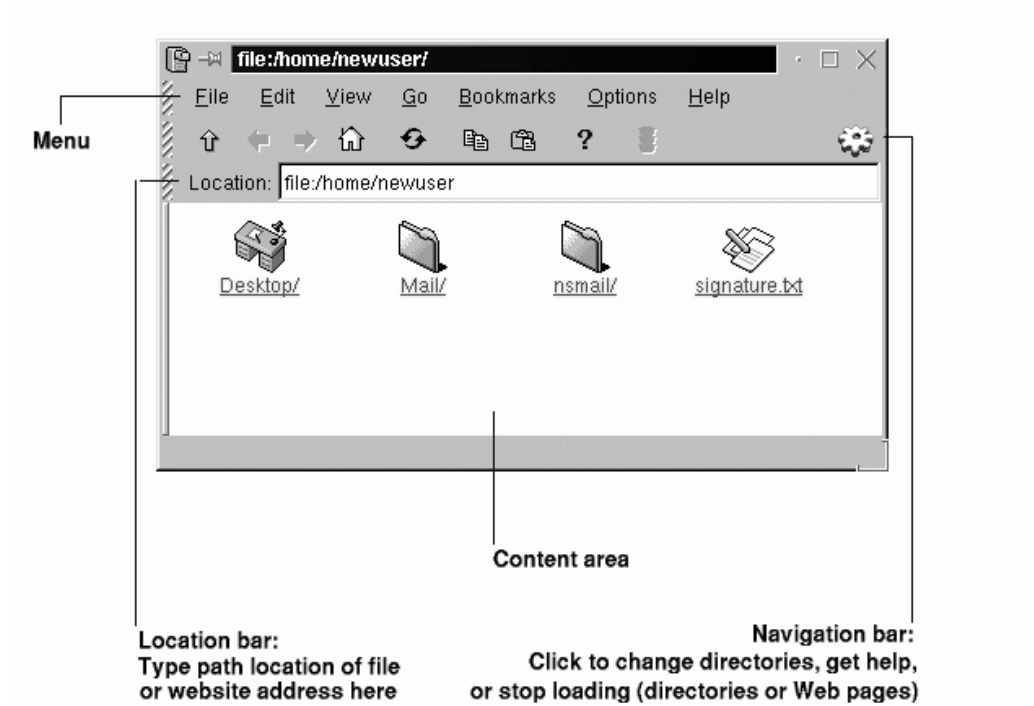
- From the **Main Menu K** icon -- Click on the **Home Directory** entry.

- From the **Panel** -- Click on the **Home Directory** icon, which looks like:




When your home directory window opens, you'll see something similar to Figure 10-6, *A Sample KFM Window*. By default, the view contains icons and links to associated files on your system. You can change this view, in the **View** menu.

Figure 10–6 A Sample KFM Window



In fact, you have quite a few choices for viewing files and directories in the content area; when you select (or deselect) a choice, the content area will **refresh** with your new view. Under **View** in the menu, you can select:

- **Show Hidden Files** -- This option shows the hidden files (also called "dot files"), mostly configuration files.
- **Show Tree** -- Gives you a **tree view**, showing you where your login directory is in the system and allowing you to navigate to other directories by clicking on other folders. Similar to Macintosh's file manager scheme, when you click on a

triangle pointed at a folder, it will **expand** a folder, revealing other folders within. Click again on the triangle, and the tree will **collapse**, and the angle will return to its initial position, like:  Root

- **Show Thumbnails** -- If you have images in a directory, selecting this option will show you tiny representations of them.
- **HTML View** -- Select this view if you want to see a directory as a Web page. This can be fun if you know HTML, or want to learn how to write a Web page. All you have to do is name your file `index.html`, and place your content as **links** in the file. When KFM opens the directory, it will open the `index.html` file.
- **Text View** and **Long View** -- Show you a directory's contents in detail, including filenames, permissions, owner, group, size and modification time. While these two views are essentially the same, small icons are included next to the listed item in **Long View** to indicate what that item represents.
- **Short View** -- Presents a list of the contents, with small icons. Doesn't present details such as owner, permissions or modification time.

Click once on a file to open the file in its associated application. For example, the file `signature.txt` is a text file. Click on it to open `kedit`, KDE's text editor.

To open a directory, click once on the file folder.

Express yourself with backgrounds

Think file managers are boring? Try livening up the content area with a colorful image background. Just right-click inside the content area of a KFM window, go to the **Dir** tab, and you can change your background. If you want to use one of the included backgrounds, select an image from the drop-down list.

Note, however, that some of the included backgrounds will probably put a strain on your eyesight -- and your sense of good taste! With names such as **spoiled_sprouts** and **ring-worm_circus**, you shouldn't expect a soothing visual experience. You can also select a background of your own by clicking on the **Browse** button and picking an image elsewhere on your system (for example, an image file in your /home directory).

If you prefer to change only the color, you can modify the color settings by going to the menu and clicking **Options => Configure File Manager**. Change the background by selecting the **Color** tab; you can change the fonts by selecting the **Font** tab.

Navigation

If you've used a Web browser or file manager before, you'll feel right at home moving around with KFM. You can type the path location of a file (and the address of a website or FTP site) in the **Location** bar. Above the **Location** bar, icons on the **Navigation bar** help you quickly move to other directories. Briefly, here's what the icons mean:

- The arrow up will take you "up" the directory tree. For example, if you're in /home/newuser and you click on the up arrow, you will move to /home; click once more, and you'll be at /, or root (the root of your filesystem, that is, not the root *account's login* directory).
-

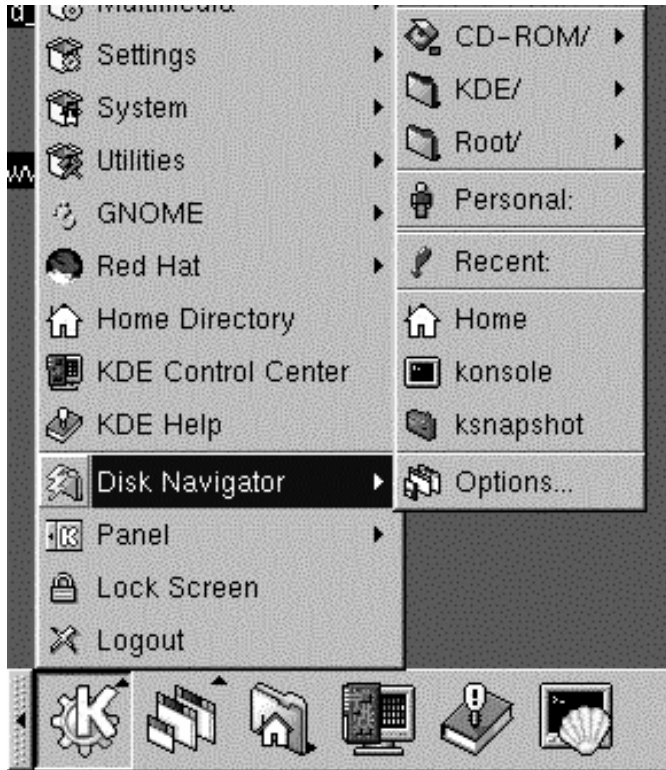
- Click on the arrow pointing *left* to take you to a previously viewed directory. Click on the arrow pointing *right* to move forward through your navigation history.
- To jump to your home location, such as `/home/newuser`, just click on the home icon.
- Click on the circular arrows to refresh your view of the current content area.
- The next two icons, of pages and a page with a clipboard, are for copy and pasting. You can copy and paste selected text in the content area, URLs or file locations from the **Location** bar, and more. If you're viewing Web pages, for example, you can copy selected content, open a text editor (**Main Menu K => Applications**) and paste the content into the editor.
- To get help, click on the question mark.
- To stop loading a page or directory in the content area, click on the traffic light.

To open a new KFM window, just click on the KDE icon (which looks like a gear) on the far right of the **Navigation bar**.

Using Disk Navigator

If you want to access a recent document or application, Disk Navigator should be the first place to turn. Using Disk Navigator, you can view directories and files on your system and return to previous files.

You can find Disk Navigator in the main menu (see Figure 10–7, *Starting Disk Navigator from the Menu*).

Figure 10–7 Starting Disk Navigator from the Menu

To work with Disk Navigator, just hold your mouse cursor over the application on the main menu; a second menu will open, showing the files and folders to which you have access. Under **Root/**, for example, you'll find directories on your system from root of the directory tree. Under the **KDE/** directory, you'll find folders that point to system documentation, games, **libraries**, and more.



Unless you know what a file is for, use caution before trying to open or modify any file.

Drag and Drop Open, Move and Copy

Keeping important documents and applications within easy reach is simple with KFM. To open a text file, for example, you can "drag" it from the content area of KFM by left-clicking on it and holding your mouse button down while you drag it. If you "drop" the file in an open text editor, it will open, ready for you to modify its contents.

In the same way, you can drag and drop files from the content area to your desktop or to other folders. When you reach your destination, a pop-up menu will offer you the choice of copying, moving or linking the file. If you select **Link** from the menu, any changes you make to a file in its new location will be reflected in the original location. A link, then, is like a shortcut to the original file.

11 An Introduction to Linuxconf

11.1 Creating a New Account in Linuxconf

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, controlling the way your system starts, and more.

Here, we'll create a user account with Linuxconf to serve as a brief introduction to the tool. For a more detailed look at the application's features -- including greater depth on working with accounts -- turn to its chapter in the system configuration section of the *Official Red Hat Linux Reference Guide*, on the Documentation CD.

Creating user accounts for yourself or others is a pretty simple process from Linuxconf, which allows you to tailor many aspects of the account.


Login as root to use Linuxconf. You can either login as root or become root in your user account. To accomplish the latter, open an Xterm and type `su`, then your root password.

Once you're root, type `linuxconf` to begin.

To learn more

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

Figure 11–1 Finding the Users Accounts Entry in 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 "+".

How to quit an Xterm window

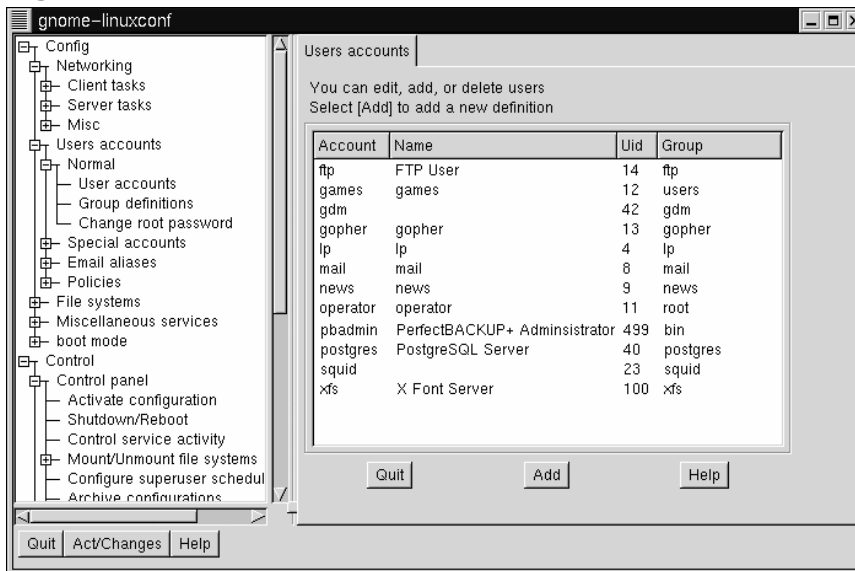
Whenever you want to quit an Xterm window, all you've got to do is type **exit** at the prompt or, in some window managers, click on the **x** in the upper right corner of the window.

Now, the menu will expand to show entries in the **Users accounts** listing. The sub-entries will look like Figure 11–2, *Users Accounts in Linuxconf*.

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 11–2, *Users Accounts in Linuxconf*).

Figure 11–2 Users Accounts in Linuxconf



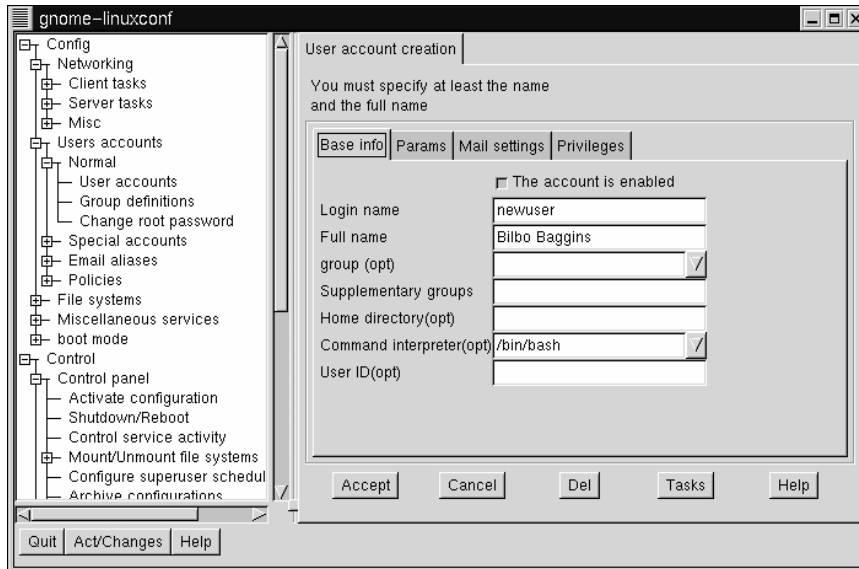
Toward the bottom of the right panel, click on the **Add** button, between the **Quit** and **Help** buttons.

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**.

Next, 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 11–3 Adding a User in Linuxconf



Accounts and groups

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

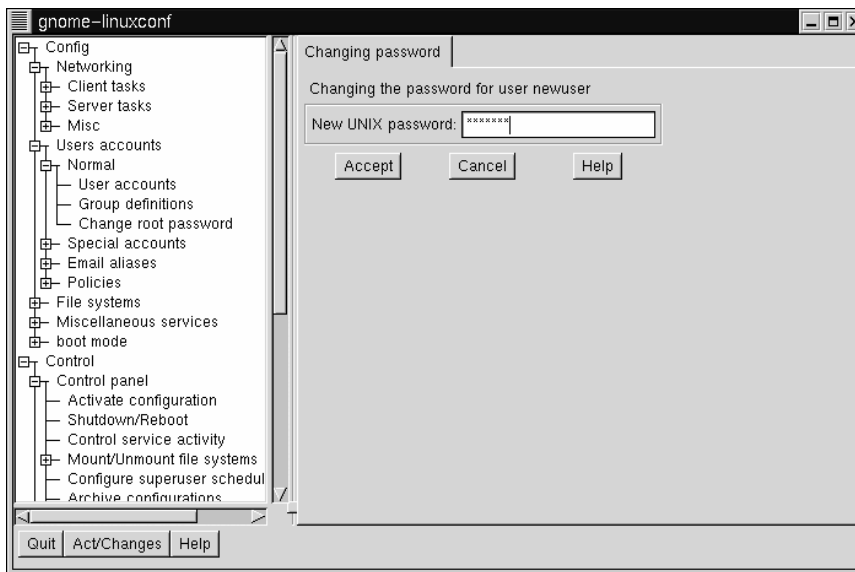
When you're finished, your entries should look like Figure 11–2, *Users Accounts in Linuxconf*.

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 use a secure password, which only you know, you've taken an important step towards keeping your system secure.

Figure 11–4 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!rp18ne, !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 11-4, *Creating a User Account Password in Linuxconf*).

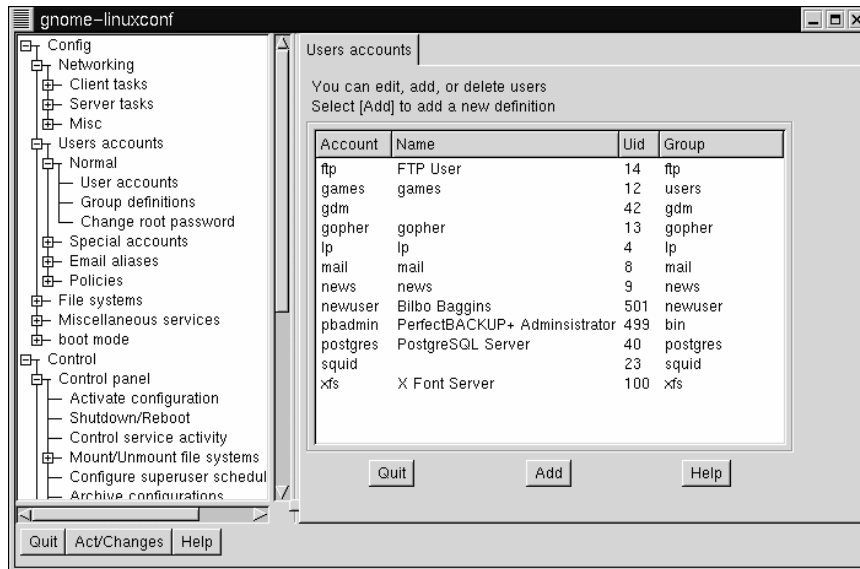
Then, click on the **Accept** button.

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 11-5, *New User Account in Linuxconf* for an example).

That's all there is to it.

Figure 11-5 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.

Is root's password secure?

Now that you've created your user account, you might want to reconsider whether your root account's password is secure 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 procedures, turn to the *System Configuration* chapter in the *Official Red Hat Linux Reference Guide* on the Documentation CD.

12 Using Gnome-RPM and RPM

12.1 The benefits of RPM

RPM stands for *RPM Package Manager*. 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 and 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* on the Documentation CD, or online at <http://www.redhat.com/support/manuals>.

You must be root

Because RPM makes changes to your system, for most functions you must be root in order to use RPM from either the shell prompt or through Gnome-RPM.

12.1.1 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` command. 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 system, 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.

12.2 Using Gnome-RPM

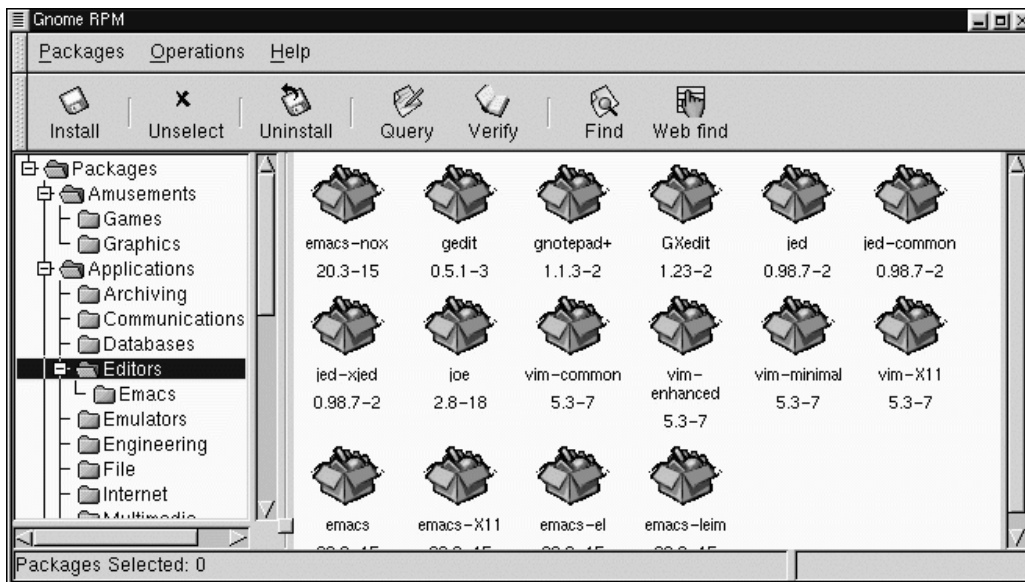
Gnome-RPM was written by James Henstridge (`james@daa.com.au`); RPM 3.0 support was written by Red Hat and additional `rpmfind` code was written by Daniel Veillard. Gnome-RPM is a fast, convenient tool for working with RPMs; with such features as **Web find**, it's also a powerful utility which allows 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 12-1, *Main Gnome-RPM Window*.

Figure 12-1 Main Gnome-RPM Window



12.2.1 The Package Window

The main window will appear similar to Figure 12-1, *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.

12.2.2 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.

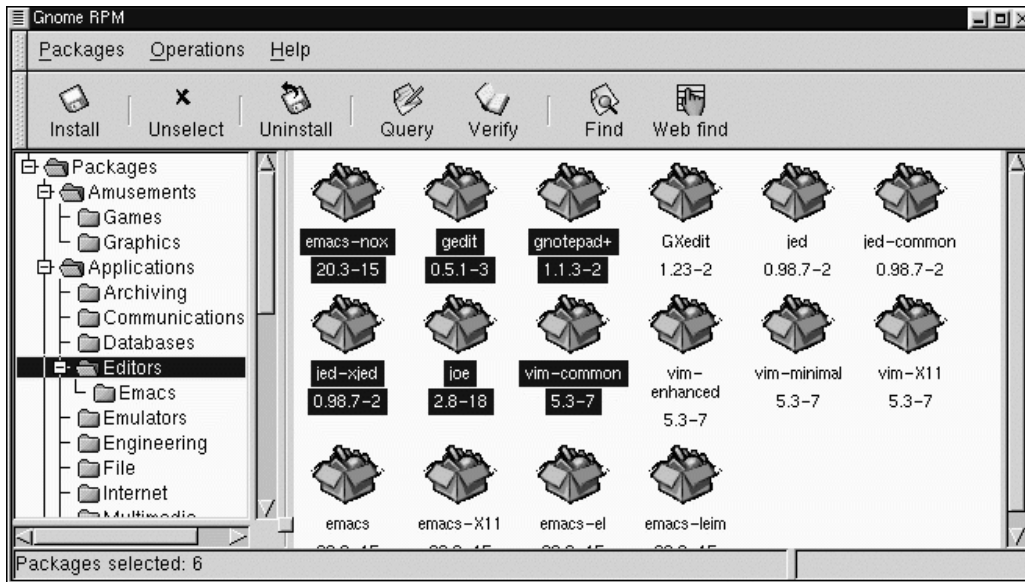
To view the packages and subgroups within a group, click once on a folder or a **+** with your left mouse button. The display window will then show you the contents of that folder.

By default, you'll see icons to represent the packages. You can customize your view by selecting **View as list** from the **Interface** tab you'll find under **Operations => Preferences**.

12.2.3 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 12-2, *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 12–2 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.

12.3 Installing New Packages

You can install new packages by choosing the **Install** button from the toolbar.

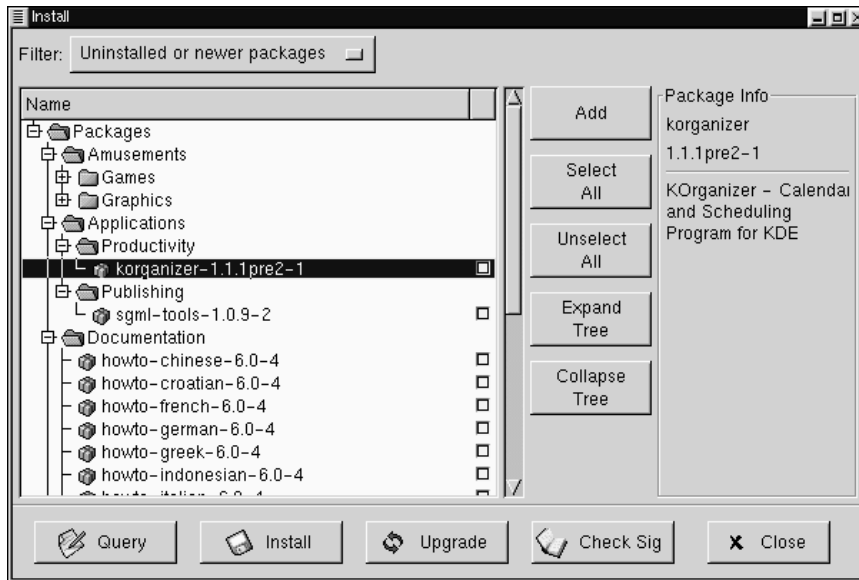
Gnome-RPM provides you with a choice of package views (as shown in Figure 12–3, *The Install Window*). You can **filter** your views for

- 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 12-3 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 12.4, *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.

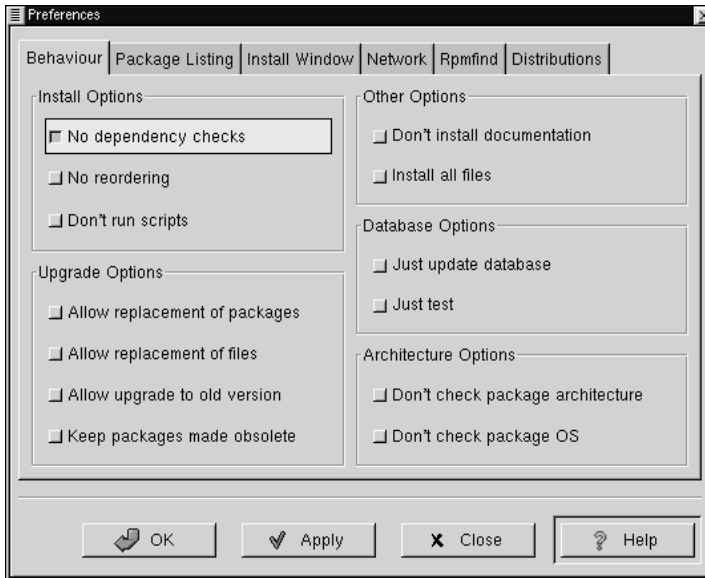
You can also "drag and drop" packages from the GNOME File Manager and place them into the **Install** window. To accomplish this, open the **File Manager (GNOME Panel => File Manager)**. Locate the directory in which your packages can be found, left-click once on the package name and "drag" the file 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.

12.4 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 12-4, *The Behavior Window* for an example.) Note that by default these boxes are not checked.

Figure 12–4 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.

Better safe than sorry

Unless you know what you're doing, you should probably avoid using this option, since packages can be dependent on a variety of programs or libraries 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.

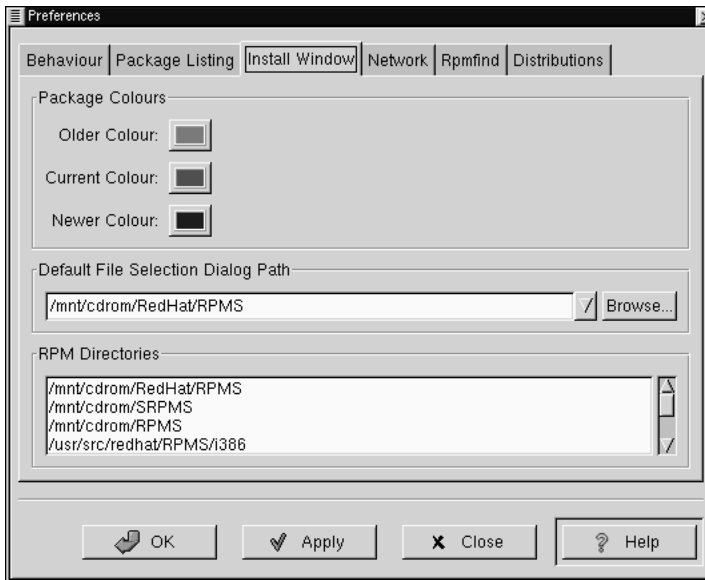
In the **Package Listing** tab, you'll find a choice of displays for your packages: either as **icons** or as a **list**, which will provide more package information than the icon display.

In **Install Window**, you can specify the path through which Gnome-RPM can find new RPMs on your system (see Figure 12–5, *The Install Window*). The default path to search 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 12–5 The Install Window



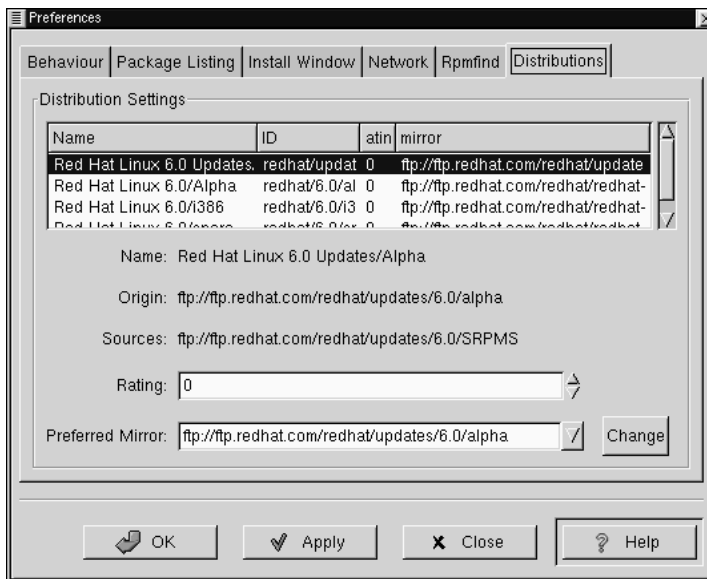
Optionally, if the path for your RPMs doesn't match the default path in your preferences, you'll be presented with a browser window, which will allow you to select the correct path for your new RPMs.

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 packages 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 12–6 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 12–6, *Distribution Settings in Preferences*) the higher the priority it will receive; the lower rating, such as "-1," will specify that packages not be recommended.

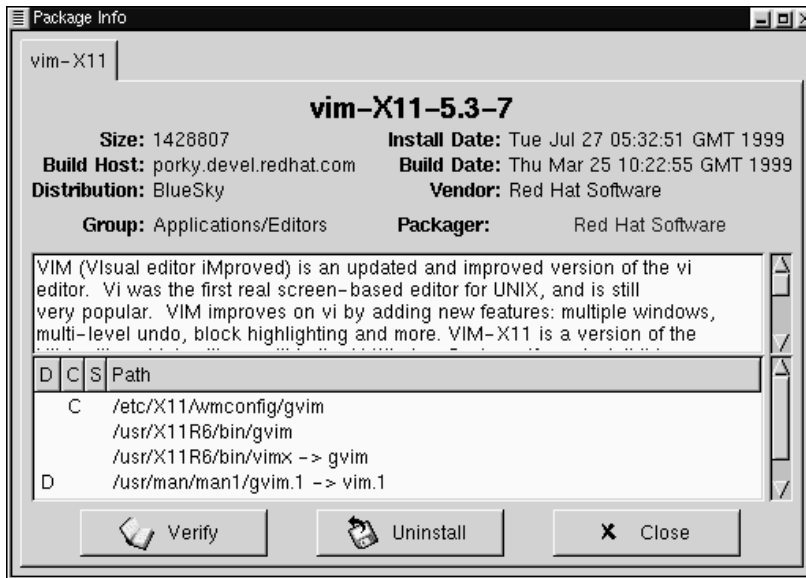
12.5 Package Manipulation

12.5.1 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 12–7, *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 12–7 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.

12.5.2 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.

12.5.3 Uninstalling Packages

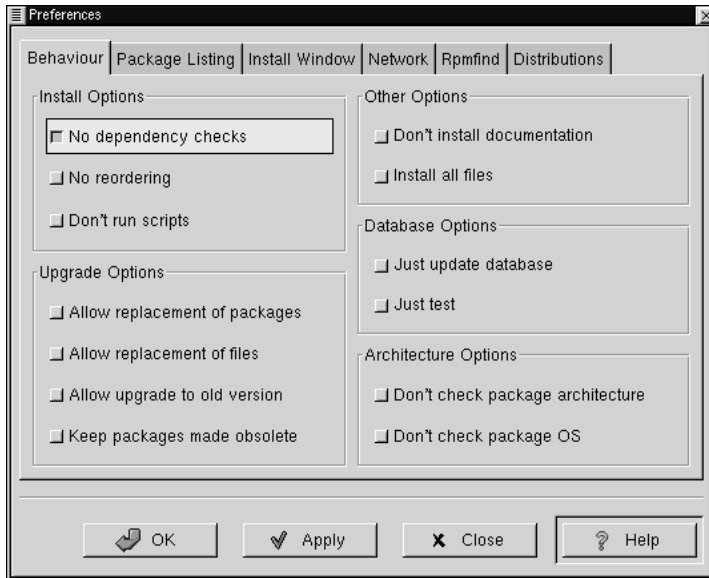
Uninstalling a package removes the application and associated files from your machine. When a package is uninstalled, any files it uses that are not needed by other packages on your system are also removed. Changed configuration files are copied to `<filename>.rpm_save` so you can reuse them later.

Make sure you're root

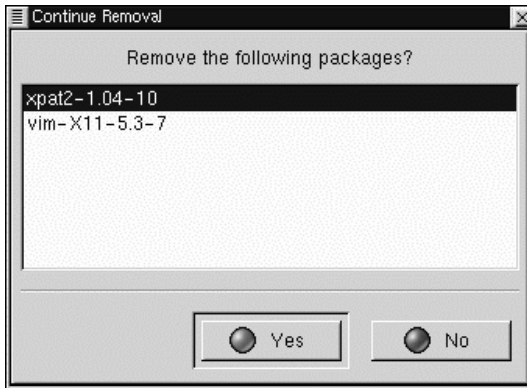
Remember that you must be root to uninstall packages.

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 12-8, *The Behavior Tab in Preferences*).

Figure 12–8 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 12–9 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 12.4, *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 12.6.2, *Uninstalling* to learn more about this process.)

12.6 RPM from the Shell Prompt

You can also use RPM technology by typing a variety 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 12.6.5, *Freshening* for more information).

For more in-depth information about RPM, refer to the related chapters in the *Official Red Hat Linux Reference Guide* on the Documentation CD. You can also read the related man page by typing `man rpm` or the help file by typing `rpm --help`.

12.6.1 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 12.6.3, *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 `--replacepks` on the command line, which tells RPM to ignore the error:

```
# rpm -ivh --replacepks 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.

12.6.2 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.

12.6.3 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:

12.6.4 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 #####
#
```

12.6.5 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.

12.6.6 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.

12.6.7 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 a 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.

12.7 Checking a Package's "Signature"

Thanks to RPM technology, it's easy to install applications instantly -- and if you don't like what you've installed, you can easily remove applications from your system.

But just because RPM is easy to use doesn't mean it's insecure. That's because RPM teams up with **Gnu Privacy Guard** -- also called **GnuPG** -- to help you make certain your downloaded package is "trustworthy."

Let's say you've just downloaded the package `coolapp-1.1.rpm`. To run a quick check of the file's integrity, you would type

```
rpm -K coolapp-1.1-1.rpm
```

Right away, you'll see the message `coolapp-1.1-1.rpm: md5 OK`. This brief message means that the file checks out -- it hasn't been corrupted by the download.

That's fine, but how trustworthy is the developer? Do you really know them? Well, if the package is **signed** with the developer's **GnuPG key**, you'll know that the developer really is who they say they are.

GnuPG is a tool for secure communication; it is a complete and free replacement for the encryption technology of **PGP**. With **GnuPG**, you can authenticate the validity of documents, encrypt and decrypt data to and from other recipients. The tool is capable of decrypting and verifying **PGP 5.x** files, as well.

During the installation of Red Hat Linux, **GnuPG** is installed by default. Here, we'll show you how to get started with **GnuPG** so you can check a developer's keys.

GnuPG uses **public key** exchange as its form of cryptography. When you check an owner's key, you can be certain that the owner is who they say they are. In public key cryptography, users create a **keypair**, which is comprised of a **public key** and a **private key**.

While you can distribute your public key to correspondents or to public **keyservers** -- online repositories which store users' public keys -- you should never divulge your private key to anyone.

12.7.1 Generating a new keypair

To use GnuPG, your first task is to generate a new keypair: your own public and private keys. You can do this with the `--gen-key` option at the shell prompt.

Do not give away your private key

Remember: While your public key can be given to anyone with whom you want to exchange secure communication, you must never give away your private key.

In an Xterm window, type `gpg --gen-key`. Since you work with your user account most frequently, you should be doing this from your user account, rather than root. You'll find an introductory screen, with key options, including one recommended option (the default), similar to the following:

```
gpg (GnuPG) 1.0.1; Copyright (C) 1999 Free Software Foundation, Inc.
This program comes with ABSOLUTELY NO WARRANTY.
This is free software, and you are welcome to redistribute it
under certain conditions. See the file COPYING for details.

Please select what kind of key you want:
  (1) DSA and ElGamal (default)
  (2) DSA (sign only)
  (4) ElGamal (sign and encrypt)
Your selection?
```

In fact, most of the screens which require you to choose an option will list the default option, within parenthesis. Unless you have a specific reason not to, you can accept the default options simply by pressing [Enter].

In the first screen, you should accept the default option: (1) DSA and ElGamal. This option will allow you to create a digital signature and encrypt (and decrypt) with two types of technologies. Type **1** and then press [Enter].

Next, choose the **key size**, or how long the key should be. Generally, the longer the key, the more resistant against attacks your messages will be. The default size, 1024 bits, should be sufficiently strong for most users, so press [Enter].

12.7.2 Exchanging Keys

Once you have created your keypair, you can exchange your public key with anyone. You do this by **exporting** your public key and **importing** others' public keys.

Exporting Keys

To export your key so that you can mail or paste it on a Web page, use the following command:

```
gpg --armor --export you@yourisp.net
```

The result will appear somewhat like this:

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: GnuPG v1.0.1 (GNU/Linux)
Comment: For info see http://www.gnupg.org

mQGiBDj8sFIRBACoyF4UucTkkuDV/KIW/kCbyzAx18OKXosVJkn4Sb/zhc/IVxoE
e/idTkzB292CUvN4KGFxCXyE0nKG/Vjc2lHwepsN41IcYQTVXUL0raITYYxyOHQe
7BJfjW0cYB/tyJLurr+iOB8JFs4HdZcCQzR9aufqcN4ErCVorrAM8k0y5wCggM9G
+Qcr32lplMKXxgf6NLj1O2oCMUwECYrrUTbxE82aZAFuJj550/QNyv0eWXWeJvzf
oPLBTaxQDp4MaYZUc5qwgxY6sxx9gBUqAJRNubg2U51K5A4emNYVRh4tRs6etpaH
+SQpA/4qRhKRHyTRK8CmGRPg9zfMkCVDHFjvA2cDXeN0Rz1Qd3WWWEGg3ytz4Omv
sYkzMAh0988HMzxWCRKXVPJoNjGHXnyvt44eb5oFE6u/b+RDcNu5FL7V7snXfOfp
GhF+YNq74tt9YF1OEjvQMRKcBXju6NnZ1pWK7fPs8W1tF2Nzo7Q1UGF1bCBHYWxs
YWdoZXIgLKhd3dy5yZWRoYXQuY29tKSA8cGF1bGdhdGxAcGVkaGF0LmNvbT6IVgQT
=Nylu
-----END PGP PUBLIC KEY BLOCK-----
```

Easily Save the Output

You can save the public key block by **redirecting** the output to a file. That way, you can insert the file in e-mail, or copy and paste it for a Web page. To redirect the output to a file called `mykey.txt` in your `/home` directory, add `> mykey.txt` to the command:

```
gpg --armor --export you@yourisp.net >
mykey.txt
```

Then the file `mykey.txt` can be inserted whenever you want to send your key to someone.

Importing Keys

When you import someone's public key, you add that key to your **keyring**. Then, when you download a document or file from that correspondent, you can check the validity of that document against the key you have added to your keyring.

To import a key, use the `--import` option. To demonstrate, let's download and import Red Hat's public key. That way, any time you want to validate a package from Red Hat, you'll be able to check it against the key you retrieved.

You can find Red Hat's key at <http://www.redhat.com/about/contact.html>. From your browser, you can download the key by pressing the [Shift] key while you click on the download link, then click the **OK** button to save the file (for example `redhat2.asc`). Now, at the shell prompt, import the key with the following command:

```
gpg --import redhat2.asc
```

The resulting message will tell you that the key was "processed." To check that the key was added, type `gpg --list-keys`. You'll see the key you'd just downloaded from Red Hat, as well as your own keys.

```
[newuser@localhost newuser]$ gpg --list-keys
/home/newuser/.gnupg/pubring.gpg
```

```
-----  
pub 1024D/DB42A60E 1999-09-23 Red Hat, Inc <security@redhat.com>  
sub 2048g/961630A2 1999-09-23
```

Keys Don't Have to be Links

Sometimes, you won't be able to download a key as a link; instead, you can save any key as a text file. As long as you know the name and location of the file you saved, you can import it to your keyring.

Verifying Packages

The command to check a downloaded package (`rpm --K filename`) is the same procedure that you used before generating your GnuPG keys and adding Red Hat's key. Now that you've added the key from Red Hat to your keyring, however, you'll see a slightly different message: `md5 gpg OK`. That means both that the package hasn't been corrupted, and that you can trust the originator of the package -- that is, Red Hat.

Going from Here

There's quite a bit more to GnuPG -- and encryption technology -- than we've covered here. In fact, there's much more information than we can cover in this book. But you can get a start on such concepts as **key management**, how to **import** and **export** keys, and more by turning to Chapter 13, *Getting Started with GnuPG*.

13 Getting Started with GnuPG

13.1 Introducing GnuPG

Have you ever wondered if your e-mail is really safe from snoops? It's easy enough to assume that your intended recipient will read your message, but unfortunately, it's also conceivable that complete strangers could intercept or even tamper with your e-mail.

In traditional "snail mail," letters are usually sealed within envelopes, stamped and delivered from post office branch to branch until they reach their destination. But sending mail through the Internet is much less secure; e-mail is usually transmitted as clear text from server to server. No special steps are taken to protect your correspondence from prying eyes or malicious behavior.

To help you protect your privacy, Red Hat Linux 6.2 includes GnuPG, the GNU Privacy Guard, which is installed by default during a typical Red Hat Linux installation.

GnuPG is a tool for secure communication; it is a complete and free replacement for the encryption technology of **PGP (Pretty Good Privacy)**, a widely popular encryption application). With GnuPG, you can encrypt your data and correspondence, and authenticate your correspondence by **digitally signing** your work. GnuPG is also capable of decrypting and verifying PGP 5.x.

Because GnuPG is compatible with other encryption standards, your secure correspondence will probably be compatible with e-mail applications on other platforms, such as Windows and MacOS.

GnuPG uses **public key cryptography** to provide users with a secure exchange of data. In a public key cryptography scheme, you generate two **keys**: a **public key** and a **private key**. You exchange your public key with correspondents or with a **keyserver**; you should never reveal your private key.

Encryption depends upon the use of keys (think of them as secret encoder/decoder rings in data format). In conventional or symmetric cryptography, both ends of the transaction have the same key, which they use to decode each other's transmissions.

In public key cryptography, two keys co-exist: a public key and a private key. A person or an organization keeps their private key a secret, and publishes their public key. Data encoded with the public key can only be decoded with the private key; data encoded with the private key can only be decoded with the public key.

More information from the shell prompt

Like most system tools for Red Hat Linux, you'll find documentation on GPG in the man pages and info pages. At a shell prompt, just type `man gpg` or `info gpg` for a quick reference of GPG commands and options.

For the most part, cryptography is beyond the scope of this publication; volumes have been written about the subject. In this chapter, however, we hope you'll gain enough understanding about GnuPG to begin using cryptography in your own correspondence. For more information about GnuPG, including an online users guide, visit <http://www.gnupg.org/>. If you want to learn more about GPG, PGP and encryption technology, turn to the references in Section 13.1.6, *Where to find more information*.

There are several graphical interfaces for GnuPG, such as SeaHorse, which is included with Red Hat Linux 7.0. However, many tasks in GnuPG can be performed quickly from the shell prompt. That's what we'll concentrate on in this chapter.

13.1.1 Generating a Keypair

To begin using GnuPG, you must first generate a new **keypair** -- that is, a public and a private key.

While instructions for generating a keypair have been previously covered (see Section 12.7.1, *Generating a new keypair*), here's a quick summary of the steps involved.

To generate a keypair, open an Xterm (or, terminal) window and, at the prompt, type

```
gpg --gen-key
```

You'll be asked to select from a number of options, such as the kind of key you want to create and the size of the key. Default options will also be available, which you

can choose by simply pressing the [Enter] key. Unless you have a reason for choosing otherwise, you'll be fine selecting the default choices.

Enter a user ID, with your name, e-mail address and an optional comment. Then, enter a **passphrase** which, like your account passwords, must be unique and secure enough to insure that no one will be able to guess it. A passphrase can be either a single word or a phrase you'll type at the prompt which identifies you as the owner of the keypair. Of course, whatever you choose, make certain that you can remember it!

A soft spot

Public key cryptography is strong, but not bulletproof. One of the weakest aspects of this technology is the vulnerability of passphrases or passwords. Picking a passphrase such as your birthdate or your significant other's address is not a good idea. Try to pick something out of the ordinary, with a mix of letters and numbers. For your passphrase, you should be as cautious as you would be with your login password.

After your public and private keys are created, they'll be placed in the `.gnupg` directory. To list your keys, type `gpg --list-keys`. You'll see your new key, similar to the following:

```
[newuser@localhost newuser]$ gpg --list-keys
/home/newuser/.gnupg/pubring.gpg
-----
pub  1024D/B7085C8A 2000-04-18 Your Name <you@yourisp.net>
sub  1024g/E12AF9C4 2000-04-18
```

13.1.2 Generating a Revocation Certificate

Once you have created your keypair, you should create a **revocation certificate** for your public key. If you should forget your passphrase, or if it's been discovered, you can publish this certificate to inform users that this public key should not be used.

Why "revoke" a key you just created?

When you generate a revocation certificate, you're not revoking the key you just created. Instead, you're giving yourself a safe way to revoke your key from public use. Let's say you create a key, then you forget your passphrase, switch ISPs (addresses), or suffer a hard drive crash. Suddenly, that revocation certificate can be a handy way to disqualify your public key.

Your signature will be valid to others who read your correspondence before your key is revoked, and you will be able to decrypt messages received prior to its revocation. To generate a revocation certificate, use the `--gen-revoke` option.

```
[newuser@localhost newuser]$ gpg --output revoke.asc
--gen-revoke <you@yourisp.net>
```

Note that if you omit the `--output revoke.asc` option from the above, your revocation certificate will be returned to the "standard output," which is your monitor screen. While you can copy and paste the contents of the output into a file of your choice using a text editor, such as **PICO**, it's probably easier to send the output to a file in your login directory. That way, you can keep the certificate handy for use later, or move it to a floppy disk and store it someplace safe.

The transaction for creating a revocation certificate will look like the following:

```
[newuser@localhost newuser]$ gpg --output revoke.asc
--gen-revoke <you@yourisp.net>

sec 1024D/823D25A9 2000-04-26 Your Name <you@yourisp.net>

Create a revocation certificate for this key? y

You need a passphrase to unlock the secret key for
user: "Your Name <you@yourisp.net>"
1024-bit DSA key, ID 823D25A9, created 2000-04-26

ASCII armored output forced.
```

```
Revocation certificate created.
```

Once your revocation certificate has been created (`revoke.asc`), it will be located in your login directory. You should make sure the file is safe, so your best bet is to copy it to a floppy diskette and store it in a secure place. (If you don't know how to copy a file to a diskette in Red Hat Linux, see Section 17.2, *Using a Diskette*.)

13.1.3 Exporting your Public Key

Before you can use public key cryptography, others must actually have a copy of your public key. To send your key either to correspondents, or to a keyserver, you must **export** the key.

To export your key, so you can display it on a Web page or paste it in e-mail, type the following:

```
[newuser@localhost newuser]$ gpg --armor --export  
<you@yourisp.net> > mykey.asc  
[newuser@localhost newuser]$
```

You won't see anything, because not only did you export your public key, you redirected the output to a file called, for example, `mykey.asc`. (Without the addition of `> mykey.asc`, the key would have been displayed as the standard output -- on the monitor screen.)

Now, the file `mykey.asc` can be inserted into e-mail or exported to a keyserver. To see the key, type `less mykey.asc` to open the file in a pager (type [Q] to quit the pager). It should look like the following:

```
-----BEGIN PGP PUBLIC KEY BLOCK-----  
Version: GnuPG v1.0.1 (GNU/Linux)  
Comment: For info see http://www.gnupg.org  
  
mQGIBDkHP3URBACKWGsYh43pkXU9wj/X1G67K8/DSr185r7dNtHnFLL/ewil10k2  
q8saWJn26QZPsDVqdUJModHfJ6kQTAt9NzQbgcVrxLYNfgeBsvkHF/PotnYcZRg  
L tZ6syBBWs8JB4xt5V09iJSGAMPUE8Jpdn2aRXPapdoDwl79LM8Rq6r+gwCg5Z  
Za pGN1kgFu24WM5wC1zg4QTbMD/3MJCSxfl99Ek5HXcb3yhj+o0LmIrGAVBgo  
WdrRd BIGjQQFhV1NSwC8YhN/4nGHWpaTxgEtnb4CI1wI/G3DK9o1YMyRJinkG  
J6XYfP3b ccQmqATDF5ugIamdditnw7dexqn/eavaMxRXJM/RQSgJJyVpbA020q  
Ke6L6Inb5H kjcZA/9obTm499dDMRQ/CNR92fA5pr0zriy/ziLUow+cqI59nt+  
bEb9nY1mfMUN6
```

```

SW0jCH+pIQH51erV+EookyOyq3ocUdjerYF/d2j19xmeSyL2H3tDvnuE6vgqFU/N
sdvby4B2Iku7S/h06W6GPQAE+pzdYX9vS+Pnf8osu7W3j60WprQkUGF1bCBHYWxs
YWdoZXIgpHBhdWxnYWxsQHJlZGhhdC5jb20+iFYEExECABYFAjkHP3UECwoEAwMV
AwIDFgIBAheAAAoJEJEcmvGCPSPmjqAoNF2zvRgdR/8or9pBhu95zeSnkb7AKCm
/uXVS0a5KoN7J61/1vEwx11poLkBDQQ5Bz+MEAQA8ztcWRJjW8cHCgLaE402jyqQ
37gDT/n4VS66nU+YItzDFScVmgMuFRzhibLb1fO9TpZzxEbSF3T6p9hLLnHCQ1bD
HRsKfh0eJYMMqB3+HyUpNeqCMEEd9AnWD9P4rQtO7Pes38sV0lX00SvsTyMG9wEB
vSNZk+Rl+phA55r1s8cAAwUEAJjqazvk0bgFrw1OGO9m7fEeDlvPSV6HSA0fvz4w
c7ckfpuxg/URQNf3TJA00Acprk8Gg8J2CtebAyR/sP5IsrK511luGdk+10M85FpT
/cen2OdJtToAF/6fGnIkeCeP105aWTbDgdAUHBrykpdWU3GJ7NS6923fVg5khQWg
uwrAiEYEGBECAAYFAjkHP4wACgkQkQKa8YI9JamliwCfXox/HjlorMKnQRJkeBcZ
iLyPH1QAoI33Ft/0HBqLtqdtP4vWYQRbibjW
=BMEc
-----END PGP PUBLIC KEY BLOCK-----

```

Exporting to a Keyserver

If you're only writing to a few correspondents, you can export your public key and send it to them personally. If there are many with whom you correspond, however, distribution of your key can be time consuming. That's where keyservers come in handy.

Figure 13–1 The Home Page of Keyserver.Net



A keyserver is a repository on the Internet which can store and distribute your public key to anyone who requests it. There are many keyservers available, and most try to remain synchronized with each other; sending your key to one keyserver is like distributing it to them all. All a correspondent has to do, then, is request your public key from a from a keyserver, import that key to their **keyring**, and they're all set (a keyring is a file in which public and secret keys are kept).

Which keyserver to use?

Because most keyservers try to be synchronized, sending your public key to one keyserver is usually as good as sending it to them all. You can, however, locate different keyservers. One place to begin your search for keyservers and more information is *Keyserver.Net*, at <http://www.keyserver.net>; another location is *Robert's Crypto & PGP Links: Keyservers*, at <http://crypto.yashy.com/www/Keyservers/>

You can send your public key from either the shell prompt or from a browser (as in Figure 13–1, *The Home Page of Keyserver.Net*); of course, you must be online to send or receive keys from a keyserver.

- From the shell prompt, type the following:

```
gpg --keyserver search.keyserver.net --send-key you@yourisp.net
```

- From your browser, go to Keyserver.Net (<http://www.keyserver.net>) and select the option to add your own PGP public key.

Your next task is to copy and paste your public key into the appropriate area on the Web page. Here's how:

- Open your exported public key file (such as *mykey.asc*, which was created in Section 13.1.3, *Exporting your Public Key*) with a pager -- for example, `less mykey.asc`
 - Using your mouse, copy the file by highlighting all the lines from the `BEGIN PGP` to `END PGP` notations (see Figure 13–2, *Copying Your Public Key*)
 - Paste the contents of the file *mykey.asc* into the appropriate area of the page on Keyserver.Net by middle-clicking with your mouse (or left- and right-clicking if you're using a two-button mouse). Then select the **Submit** button on the keyserver page. (If you make a mistake, press the **Reset** button on the page to clear your pasted key.)
-

Figure 13–2 Copying Your Public Key

```

paulgall@pinky.meridian.redhat.com: /home/saqadoc/paulgall
File Edit Settings Help
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: GnuPG v1.0.1 (GNU/Linux)
Comment: For info see http://www.gnupg.org

mQ6iBDkHP3URBACKU6sYh43pkXU9wj/X1667K8/DSr185r7dNtHNFLL/ewil10kZ
q8saUjN26QZPsDUvdUJM0dHfJ6kQTAt9NzQbgcVrxLYNfgeBsvkHF/P0tnYcZRgL
tZ6syBBWs8JB4xt5U09iJS6AMPUQE8Jpdn2aRXPapdoDw179LM8Rq6r+gwCg5ZZa
p6N1kgFu24W5wC1zg4QTbMD/3MJCSxfL99Ek5HXcB3yhj+o0LmIr6AUVBgoWdrRd
BIGjQQFhV1NSwC8YhN/4nGHwpaTxgEtnb4CI1wI/63DK9o1YMyRJink6J6XYfP3b
cCQmqATDF5ugIamdditn7deXqn/eavaMxRXJM/RQsGJJyUpbA020qKe6L6Inb5H
kjcZA/9obTm499dDMRQ/CHR92fA5pr0zriy/ziLUow+cqI59nt+bEb9nY1mfMUN6
SW0jCH+pIQH51erV+Eooky0yq3ocUdjeRYF/d2j19xmeSyl2H3tDvnuE6vgqFU/N
sdvby4B2Iku7S/h06W66PQae+pzdyX9vS+Pnf8osu7W3j60WprQkUGF1bCBHYWxs
YWdoZXI9PHBhdWxnYWsQHJlZ6hhdC5jb20+iFYEEcABYFAjKHP3UECwoEAwMU
AwIDFgIBaheAAaAJEJECmV6CPSVpMjQAoHF2zvRgdR/8or9pBhu95zeSnkb7AKCn
/uXUS0a5KoN7J61/1vEwx11poLkBDQQ5Bz+MEaQA8ztcWRJjW8cHCgLaE402jyqQ
37gDT/n4US66nU+YItzDFScUmgMuFRzhibLb1f09TpZzxEbSF3T6p9hLLnHCQ1bD
HRsKfh0eJYMMqB3+HyUpNeqCMEEd9AnWD9P4rQt07Pes38sU01X00SvsTyMG9wEB
vSNZk+R1+phA55r1s8cAAwUEAJjqazvk0bgFrw10609m7fEeD1vPSU6HSA0fvz4w
c7ckfpuxg/URQNf3TJA00AcpRk86g8J2CtebAyR/sP5IsrK511luGdk+10M85FpT
/cen20dJtToAF/6f6nIkeCeP105aWTbDgdAUHBRykpDWU36J7NS6923fVg5khQWg
uwrAiEYEGBECAAYFAjKHP4wACgkQKqKa8YI9JamliwCfXox/HjlorMKnQRJkeBcZ
iLyPH1QAoI33Ft/0HBqLtqdtP4vWYQRb1bjW
=BMEc
-----END PGP PUBLIC KEY BLOCK-----
~
(END)

```

Note that if you are submitting your key to another Web-based keyserver, the above transaction will be essentially the same.

That's it. Regardless of whether you use the shell prompt or the Web, you'll see a message that your key was successfully submitted -- either in the Xterm window or at the keyserver's website. Now, users who want to communicate securely with you can **import** your public key and add it to their keyring.

13.1.4 Importing a Public Key

The other end of key exchange -- importing others' public keys to your keyring -- is just as simple as exporting keys. When you import someone's public key, you can decrypt their mail and check their digital signature against their public key on your keyring.

One of the easiest ways to import a key is to download the key or save it from a website (as in saving Red Hat's key, for example, in *Importing Keys* in Section 12.7.2).

After downloading a key, use the `--import` command to add it to your keyring.

Another way to save a key is to use a browser's **Save As** feature. If you're using a browser such as **Navigator**, and you locate a key at a keyserver, you can save the page as a text file (go to **File** => **Save As**). In the drop-down box next to **Format for saved document**, choose **Text**. Then, you can import it -- as long as you remember the name of the file you saved. For example, let's say you just saved a key as a text file, called *newkey.txt*. To import the file, at a shell prompt, type:

```
[newuser@localhost newuser]$ gpg --import newkey.txt
gpg: key F78FFE84: public key imported
gpg: Total number processed: 1
gpg:          imported: 1
[newuser@localhost newuser]$
```

To check that the process was successful, use the `--list-keys` command at the prompt, as in `gpg --list-keys`; you should see your newly imported key listed on your keyring.

13.1.5 What Are Digital Signatures?

Digital signatures are similar to signatures -- or stamps -- you might place on a document such as e-mail.

Unlike traditional correspondence, in which it might be possible to tamper with your written signature, digital signatures can't be forged. That's because the signature is created with your unique secret key, and can be verified by your recipient using your public key.

A digital signature **timestamps** a document; in essence, that means that the time you signed the document is part of that signature. So if anyone tries to modify the document, the verification of the signature will fail. Some e-mail applications, such as Exmh or KDE's KMail, include the ability to sign documents with GnuPG within the application's interface.

Two useful types of digital signatures are **clearsigned** documents and **detached signatures**. Both types of signatures incorporate the same security of authenticity, without requiring your recipient to decrypt your entire message.

In a clearsighted message, your signature appears as a text block within the context of your letter; a detached signature is sent as a separate file with your correspondence (that is, detached from your message).

13.1.6 Where to find more information

There's more to encryption technology than can be covered in one slim introduction to GPG. Here are some resources where you can learn more:

- Websites:
 - **GnuPG:** Go to <http://www.gnupg.org> for links to the latest GPG releases, a comprehensive user's guide, and other cryptography resources.
 - *An Encryption Tutorial*, from Webmonkey: Go to <http://hotwired.lycos.com/webmonkey/backend/security/tutorials/tutorial1.html> to learn more about encryption and how to apply encryption techniques.
 - The Electronic Frontier Foundation, "Privacy, Security, Crypto, & Surveillance" Archive: <http://www.eff.org/pub/Privacy>.
 - Books:
 - *The Official PGP User's Guide*, Philip R. Zimmermann; MIT Press, 1995
 - *PGP: Pretty Good Privacy*, Simson Garfinkel; O'Reilly & Associates, 1994
 - *E-Mail Security: How to Keep Your Electronic Messages Private*, Bruce Schneier; John Wiley & Sons, 1995
-

Part III The Shell Prompt

14 Shell Prompt Basics

14.1 Why Use a Shell Prompt?

Graphical environments for Linux have come a long way in the past few years. You can be as productive as you'd like in the X Window System, without ever having to open a shell prompt.

So why spend the time learning how to accomplish tasks from the shell prompt? For one thing, you can work faster. In the time it might take you to open the file manager in GNOME or KDE, then locate a directory, then create, delete or modify files, you could have finished your work with several commands from a shell prompt.

In this section, we'll show you how you can move around (without getting lost), manipulate files, perform some basic administration tasks and more. All the work can be performed from within an Xterm window, too, so it won't be necessary to forsake the GUI to work faster from the shell prompt.

14.2 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.

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.

Open an Xterm and make sure you're logged in as root. At the shell prompt, if you see something like `[newuser@localhost newuser]$`, for example, type:

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

A "super" command

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

Put a standard diskette in the floppy drive.

Naming the floppy drive

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

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

At the prompt, type:

```
uname -r
```

Your kernel version will be displayed. 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, for example: 2.2.14-5.0).

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].

Clearing things up

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.

14.3 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 in a graphical presentation:

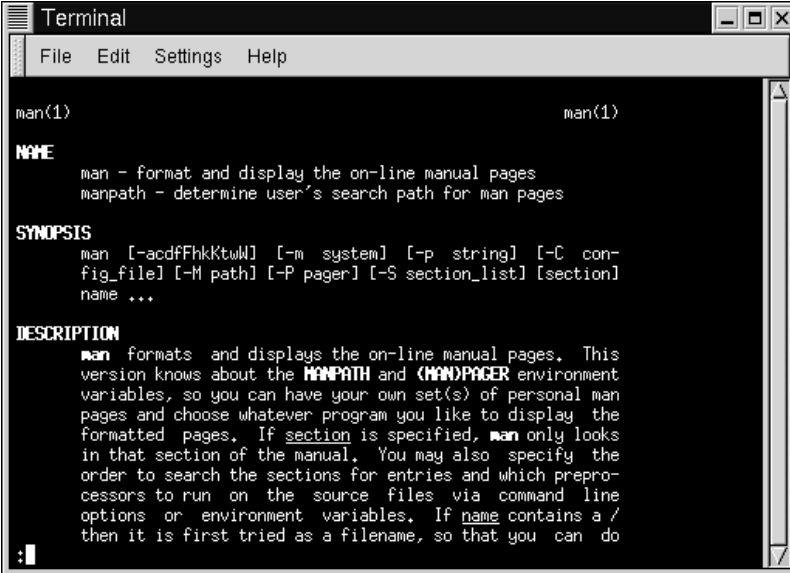
- from the GNOME help browser (see Section 2.6, *Finding Help*)
- from the KDE help browser (see Section 3.7, *Finding Help*)
- from an application called `xman` at the shell prompt

You can quickly access a man page for a particular page by typing the following at the shell prompt:

```
man command
```

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

Figure 14–1 Reading a Man Page at the Shell Prompt

A terminal window titled "Terminal" with a menu bar (File, Edit, Settings, Help) and window controls. The terminal shows the output of the 'man(1)' command. The output is formatted with bold headers: NAME, SYNOPSIS, and DESCRIPTION. The NAME section lists 'man' and 'manpath'. The SYNOPSIS section shows the command syntax with options like [-acdfFhkKtW] and [-m system]. The DESCRIPTION section explains that 'man' formats and displays on-line manual pages, mentioning environment variables like MANPATH and MANPAGER.

```
man(1) man(1)
NAME
man - format and display the on-line manual pages
manpath - determine user's search path for man pages

SYNOPSIS
man [-acdfFhkKtW] [-m system] [-p string] [-C con-
fig_file] [-M path] [-P pager] [-S section_list] [section]
name ...

DESCRIPTION
man formats and displays the on-line manual pages. This
version knows about the MANPATH and MANPAGER environment
variables, so you can have your own set(s) of personal man
pages and choose whatever program you like to display the
formatted pages. If section is specified, man only looks
in that section of the manual. You may also specify the
order to search the sections for entries and which prepro-
cessors to run on the source files via command line
options or environment variables. If name contains a /
then it is first tried as a filename, so that you can do
```

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 14–1, *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 **pip**ing, and you can learn more about it in Chapter 15, *You Are Here*.

15 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.

Make sure you're in 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 Section 1.1.1, *Create a User Account* if you require assistance on how to create your user account.


15.1 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 15–1 The Command `pwd` Shows You Where You Are

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

```
[newuser@localhost newuser]$
```

Now type:

```
pwd
```

What do you see? Something like

```
/home/newuser
```

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.)

15.2 Getting from Here to There: `cd`

Whenever you want to change to your home directory, all you've got to do is type:

```
cd
```

Go ahead, try it in an Xterm window.

If you were already in your home directory, that didn't do much. If you want to go somewhere other than your home directory, you'll need to tell your system where you want 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:

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

Now that you see where you are, you can 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 15–2 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/newuser`, you were looking at the full path -- or the absolute path from the root directory. You can think of the `newuser` 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:

```
[newuser@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/newuser
```

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 15–3, *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:

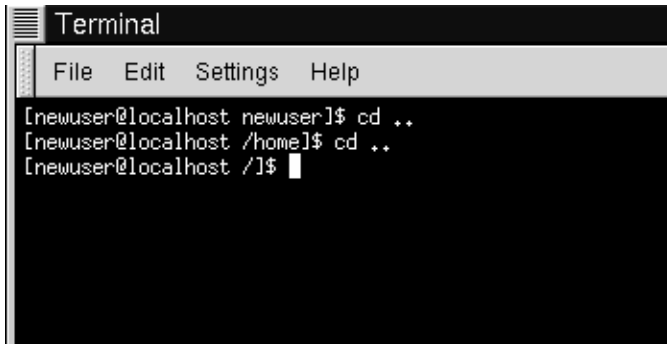
```
cd home/newuser
```

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 `newuser`, 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 15–3 Relative Pathnames are Relative to Current PositionA terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal shows three lines of command history: "[newuser@localhost newuser]\$ cd ..", "[newuser@localhost /home]\$ cd ..", and "[newuser@localhost /]\$". The cursor is at the end of the third line.

```
Terminal
File Edit Settings Help
[newuser@localhost newuser]$ cd ..
[newuser@localhost /home]$ cd ..
[newuser@localhost /]$
```

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

What are parent directories?

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 `newuser`.

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 shopping 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.

Is the path absolute or relative?

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/newuser
```

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 `newuser`, in the `home` directory."

Get back to your login directory quickly

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.

Know where you are

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.

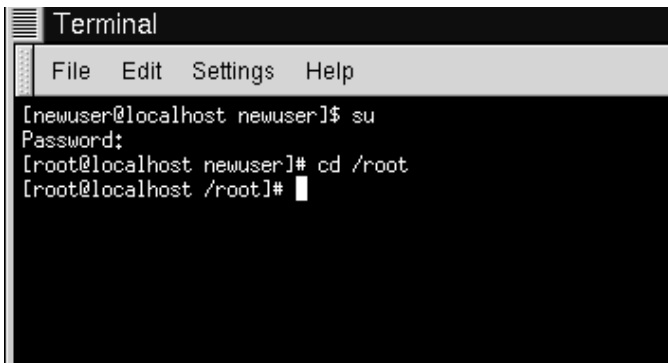
Denying access to the root and other users' accounts (or **login directories**) is one way your Linux system prevents accidental or malicious tampering. You'll find out more about file "ownership" and permissions later in this chapter.

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

```
[newuser@localhost newuser]$ su
Password:your root password
[root@localhost newuser]#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 15–4, *Becoming Root*).

Figure 15–4 Becoming Root



```
Terminal
File Edit Settings Help
[newuser@localhost newuser]$ su
Password:
[root@localhost newuser]# cd /root
[root@localhost /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
[newuser@localhost newuser]$
```

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.

15.3 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 directory 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 of 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 16, *Working with 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:

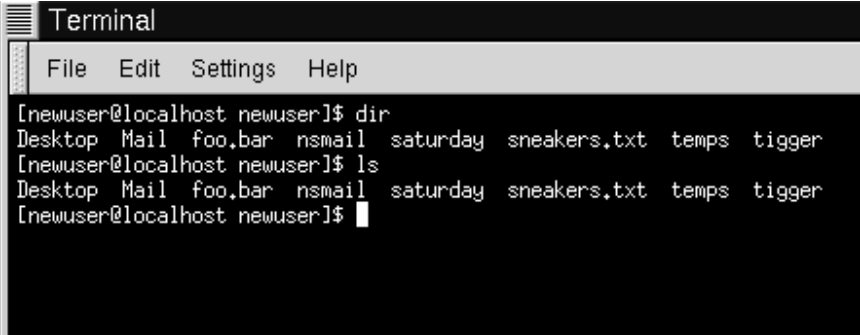
```
dir
```

Now, in the same Xterm window, type:

```
ls
```

Looks the same (see Figure 15-5, *dir and ls Seem Similar*). You see, among other contents, your new file, `foo.bar` and the new directory, `tigger`.

Figure 15-5 `dir` and `ls` Seem Similar



```
Terminal
File Edit Settings Help
[newuser@localhost newuser]$ dir
Desktop Mail foo.bar nsmail saturday sneakers.txt temps tigger
[newuser@localhost newuser]$ ls
Desktop Mail foo.bar nsmail saturday sneakers.txt temps tigger
[newuser@localhost newuser]$
```

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
```

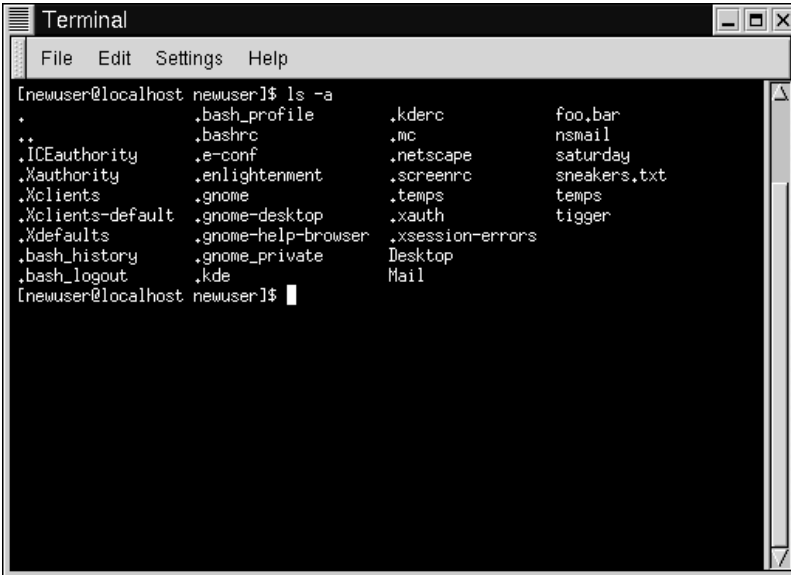
Quite a difference. When you added the `-a` option, you were specifying that you wanted to list *all* the files in the directory (see Figure 15–6, *ls with -a Option*).

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

Read the `ls` man page

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 15–6 `ls` with `-a` OptionA terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal shows the command `[newuser@localhost newuser]$ ls -a` and its output. The output lists various files and directories, including hidden files starting with a dot. The files are listed in four columns:

.	.bash_profile	.kderc	foo.bar
..	.bashrc	.mc	nsmail
.ICEauthority	.e-conf	.netscape	saturday
.Xauthority	.enlightenment	.screenrc	sneakers.txt
.Xclients	.gnome	.temps	temps
.Xclients-default	.gnome-desktop	.xauth	tigger
.Xdefaults	.gnome-help-browser	.xsession-errors	
.bash_history	.gnome_private	Desktop	
.bash_logout	.kde	Mail	

The terminal prompt `[newuser@localhost newuser]$` is visible at the end of the output.

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, ownership, 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.

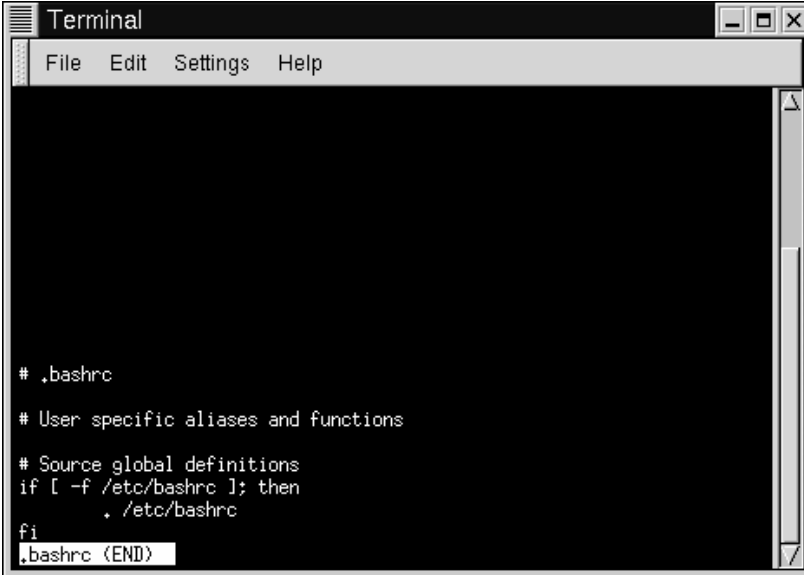
```
ls -al --color /etc
```

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 15-7, *The .bashrc File*).

Now before we go any further...

Figure 15–7 The `.bashrc` FileA terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal displays the contents of the `.bashrc` file. The text shown is:

```
# .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 simple **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; they are put there to help guide anyone who's editing or modifying files.

Move 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

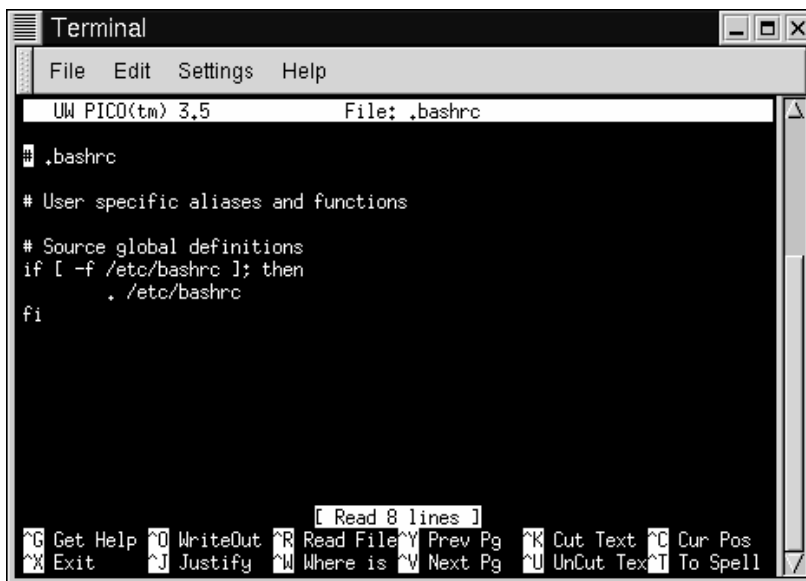
# User specific aliases and functions

alias ls="ls -al --color"

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

See Figure 15–8, *Adding an Alias for `ls` to `.bashrc`* for an example in Pico.

Figure 15–8 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 (`.filename`). 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 categorized by color.

15.4 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 (/).

What is the FHS?

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*, on the Documentation CD. You can also visit the FHS website: <http://www.pathname.com/fhs>.

As long as you're logged into our user account -- which will help prevent disastrous mistakes -- why not take a look around.

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

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

```
cd /
```

We'll see a prompt which looks like:

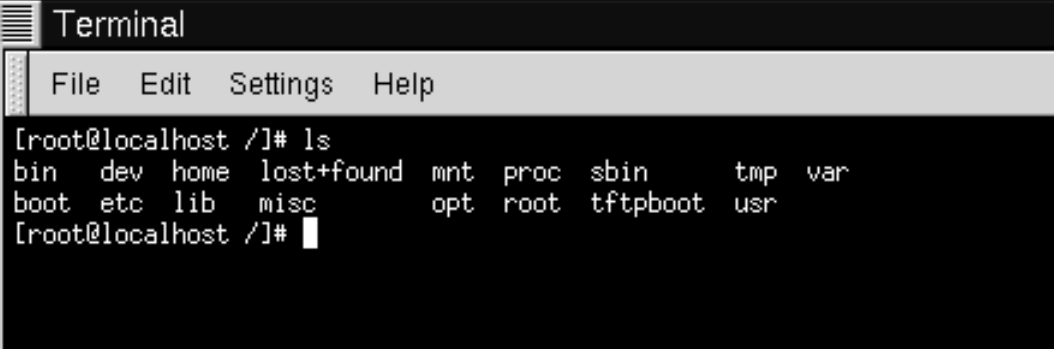
```
[newuser@localhost ~]$
```

Now, take a look at 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 15–9 A View of the Directories from Root

A terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal shows the command prompt "[root@localhost ~]#" followed by the command "ls". The output is a list of directories: "bin", "dev", "home", "lost+found", "mnt", "proc", "sbin", "tmp", "var", "boot", "etc", "lib", "misc", "opt", "root", "tftpboot", and "usr". The prompt returns to "[root@localhost ~]#" with a cursor.

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.

```
[newuser@localhost ~]$ cd etc
[newuser@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.

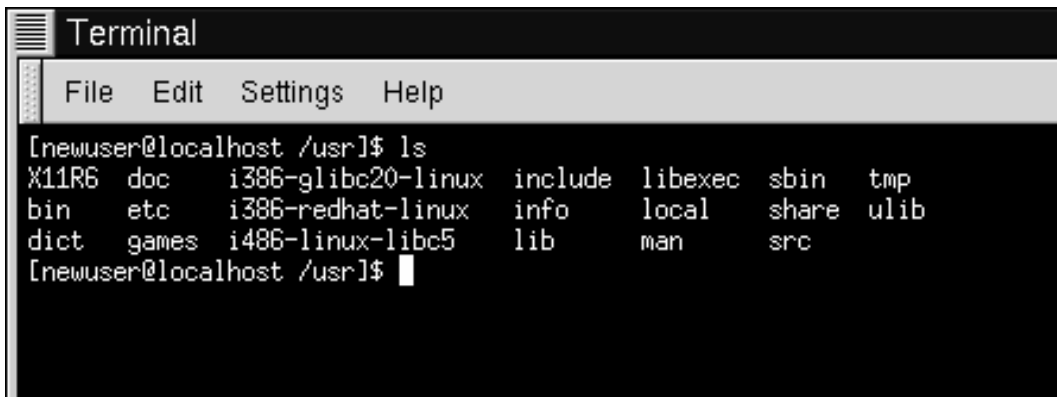
When that user account was created, files were taken from `/etc/skel` and placed into the new account. These files helped to "flesh out" the account (skeleton... flesh... get it?).

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

```
[newuser@localhost /skel]$ cd /usr
[newuser@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 15–10, *Typing `ls` in `/usr`*).

Figure 15–10 Typing `ls` in `/usr`

A terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal shows the user at the prompt [newuser@localhost /usr]\$ typing the command ls. The output is a list of directories: X11R6, bin, dict, doc, etc, games, i386-glibc20-linux, i386-redhat-linux, i486-linux-libc5, include, info, lib, libexec, local, man, sbin, share, tmp, and ulib. The prompt returns to [newuser@localhost /usr]\$ with a cursor.

```
Terminal
File Edit Settings Help
[newuser@localhost /usr]$ ls
X11R6  doc      i386-glibc20-linux  include  libexec  sbin  tmp
bin    etc      i386-redhat-linux   info     local    share ulib
dict   games   i486-linux-libc5    lib      man      src
[newuser@localhost /usr]$
```

In `/usr/share/man` we'll find the system manual pages; other documentation which isn't covered by man pages will be found in `/usr/share/doc` and in `/usr/share/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 *RPM Package Manager*) 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 12, *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`.

15.5 "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`.

15.6 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:

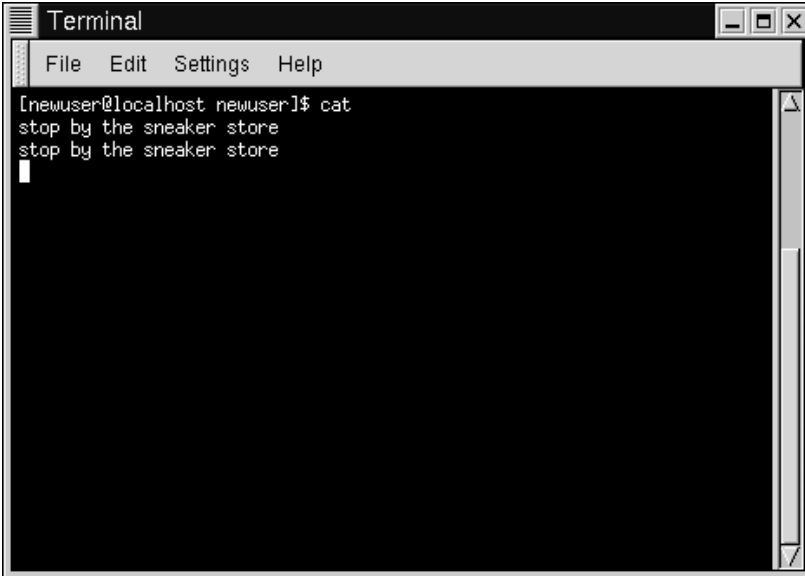
```
cat
```

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:

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

Figure 15–11 `cat` Demonstrates Standard Input and Standard OutputA screenshot of a terminal window titled "Terminal". The window has a menu bar with "File", "Edit", "Settings", and "Help". The terminal prompt is "[newuser@localhost newuser]\$", followed by the command "cat". The output of the command is "stop by the sneaker store" on two lines. The cursor is on a blank line below the output.

```
Terminal
File Edit Settings Help
[newuser@localhost newuser]$ cat
stop by the sneaker store
stop by the sneaker store
█
```

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 15–11, *cat Demonstrates Standard Input and Standard Output*).

15.7 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.

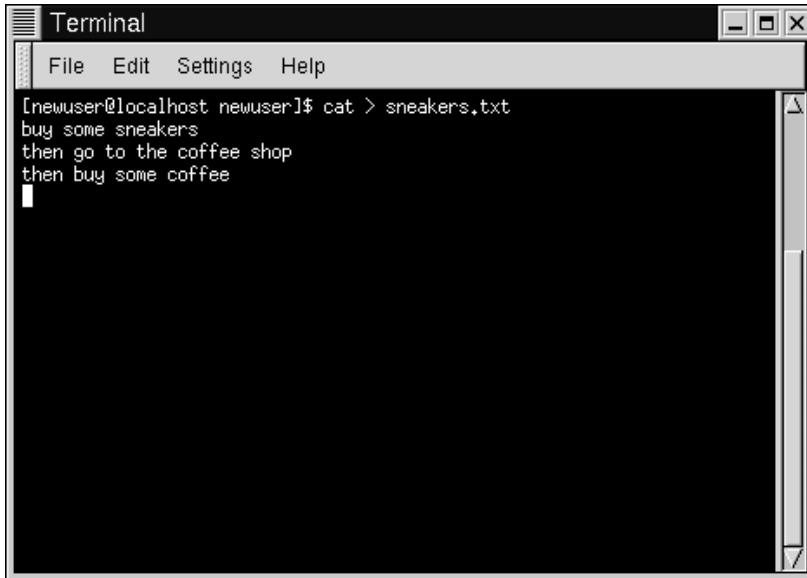
We used `cat` before to demonstrate the idea behind standard input and standard output. Now, let's use `cat` to see how standard output can be redirected.

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:

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

Figure 15–12 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 15–12, *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.



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`.

```
[newuser@localhost newuser]$ 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
```

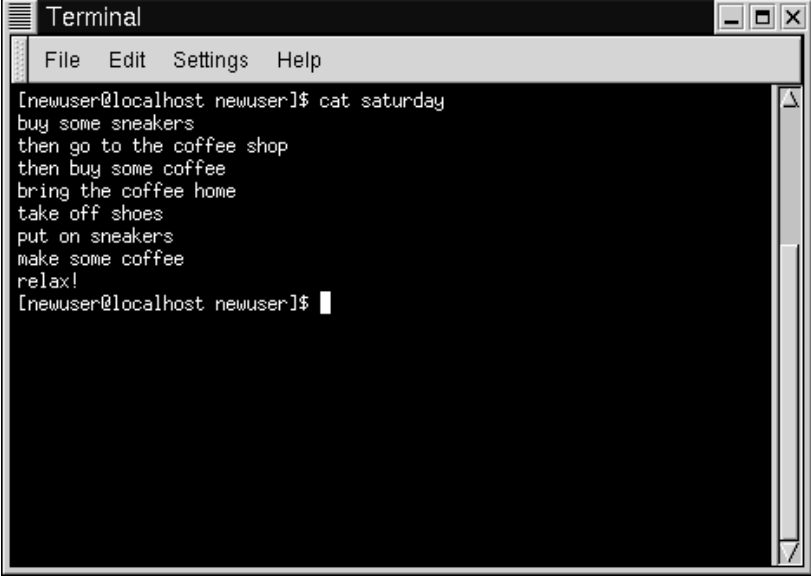
at the prompt.

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 15–13, *Joining Files and Redirecting Output*).

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

That's it.

Figure 15–13 Joining Files and Redirecting Output

A terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal shows the command `[newuser@localhost newuser]$ cat saturday` and its output: `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`, and `relax!`. The prompt `[newuser@localhost newuser]$` is visible at the bottom.

```
[newuser@localhost newuser]$ 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!
[newuser@localhost newuser]$
```

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

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

and you should see something like this:

```
[newuser @localhost newuser]$ 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!
[newuser @localhost newuser]$
```

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

Combining files with `cat`

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.

15.8 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.

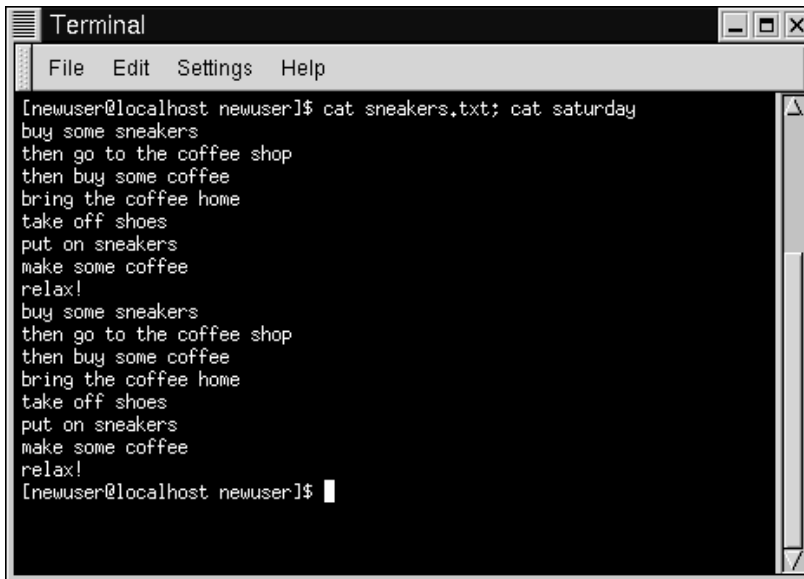
Compare the results of the files `sneakers.txt` and `saturday` now, and you'll see that they're identical. To make your comparison, just type:

```
cat sneakers.txt; cat saturday
```

The contents of both files will be displayed - first `sneakers.txt`, then `saturday` (as shown in Figure 15-14, *Stringing Commands and Comparing Files*).



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 15–14 Stringing Commands and Comparing FilesA terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal shows a shell prompt "[newuser@localhost newuser]\$" followed by the command "cat sneakers.txt; cat saturday". The output consists of two identical blocks of text, one from "sneakers.txt" and one from "saturday". Each block contains the following lines: "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", and "relax!". The terminal ends with the prompt "[newuser@localhost newuser]\$" and a cursor.

```
[newuser@localhost newuser]$ 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!
[newuser@localhost newuser]$
```

(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.`

15.9 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`.

15.10 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 15–15, *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]`.

How to read your startup messages

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

Figure 15–15 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.

15.11 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 15–14, *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:

```
[newuser@localhost newuser]$ date; netscape; date  
Mon Feb 7 13:26:27 EST 2000
```

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

```
[newuser@localhost newuser]$ date; netscape; date  
Mon Feb 7 13:26:27 EST 2000  
Mon Feb 7 14:28:32 EST 2000  
[newuser@localhost newuser]$
```

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.

15.12 Ownership and Permissions

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

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

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.

```
[newuser@localhost newuser]$ su
Password: your root password
[root@localhost newuser]# 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 15-16, *Permissions for sneakers.txt*).

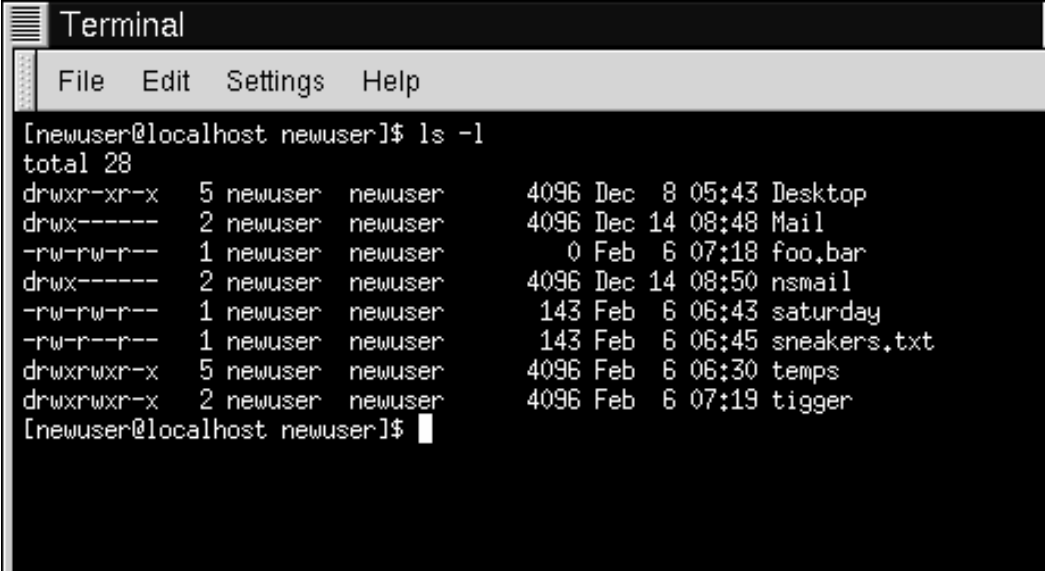
```
[newuser@localhost newuser]$ ls -l sneakers.txt
-rw-rw-r--  1 newuser newuser  150 Mar 19 08:08 sneakers.txt
```

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 (newuser) and to which group the owner belongs (newuser).

Your default group

Remember that, by default, your group was the login name you chose.

Figure 15–16 Permissions for `sneakers.txt`



```
Terminal
File Edit Settings Help
[newuser@localhost newuser]$ ls -l
total 28
drwxr-xr-x  5 newuser  newuser  4096 Dec  8 05:43 Desktop
drwx-----  2 newuser  newuser  4096 Dec 14 08:48 Mail
-rw-rw-r--  1 newuser  newuser    0 Feb  6 07:18 foo,bar
drwx-----  2 newuser  newuser  4096 Dec 14 08:50 nsmail
-rw-rw-r--  1 newuser  newuser  143 Feb  6 06:43 saturday
-rw-r--r--  1 newuser  newuser  143 Feb  6 06:45 sneakers.txt
drwxrwxr-x  5 newuser  newuser  4096 Feb  6 06:30 temps
drwxrwxr-x  2 newuser  newuser  4096 Feb  6 07:19 tigger
[newuser@localhost newuser]$
```

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 remaining 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 (newuser) and those in newuser's group (which is also newuser).

Let's stretch out these file settings a bit:

```

-      (rw-)  (rw-)  (r--)   1 newuser newuser
|        |      |      |
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.

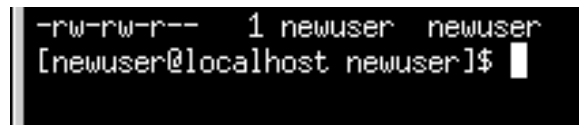
Let's look again at first column of `sneakers.txt` and identify its permissions. (See Figure 15–17, *A Closer View of Permissions*)

```

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

```

Figure 15–17 A Closer View of Permissions



```

-rw-rw-r-- 1 newuser newuser
[newuser@localhost newuser]$ █

```

The file's owner, `newuser`, has permission to read and write to the file; it's not a program, so `newuser` doesn't have permission to execute it. The group, `newuser`, has permission to read and write to `sneakers.txt`, as well. Similar to the program notation for owner `newuser`, there's no execute permission for group `newuser`.

In the last set, we can see that those who aren't either the user `newuser` or in the group called `newuser` 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 newuser newuser  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 `newuser`) can read and write to the file.

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



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 "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 newuser newuser  150 Mar 19 08:08 sneakers.txt
```

Now, type the following:

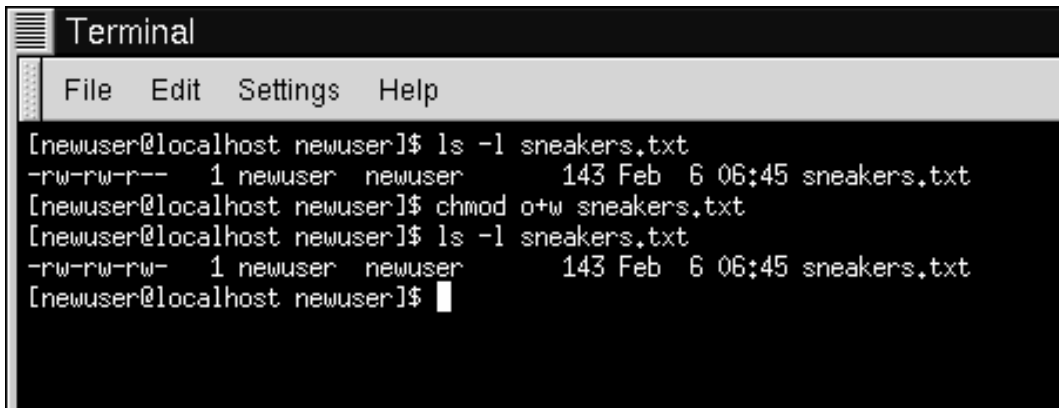
```
chmod o+w sneakers.txt
```

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

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

There's our result: Now, everyone can read and write to the file (Figure 15–18, *Changing Permissions for sneakers.txt*).

Figure 15–18 Changing Permissions for `sneakers.txt`

A terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal shows the following commands and output:

```
[newuser@localhost newuser]$ ls -l sneakers.txt
-rw-rw-r--  1 newuser newuser  143 Feb  6 06:45 sneakers.txt
[newuser@localhost newuser]$ chmod o+w sneakers.txt
[newuser@localhost newuser]$ ls -l sneakers.txt
-rw-rw-rw-  1 newuser newuser  143 Feb  6 06:45 sneakers.txt
[newuser@localhost newuser]$
```

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

If we want to remove read and write permissions 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 newuser newuser  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:

15.12 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)

15.12 Permissions

- r -- read access
- w -- write access
- x -- execute access

15.12 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 everyone.

```
chmod a-rw sneakers.txt
```

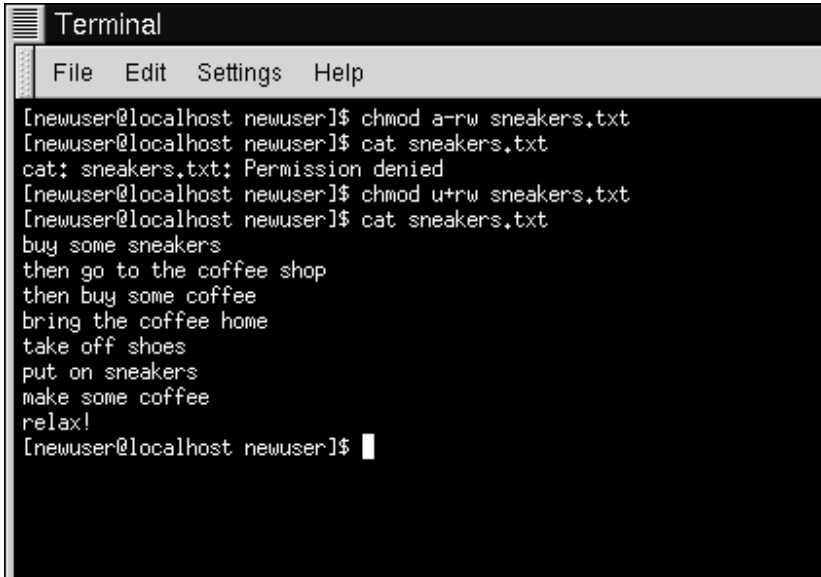
Now, let's see if we can read the file:

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

Guess it worked; even we can't get into the file. But since the file belongs to us, we can always change its permissions to allow us read and write access. (See Figure 15–19, *Removing and Restoring Permissions*)

```
[newuser@localhost newuser]$ chmod u+rw sneakers.txt
[newuser@localhost newuser]$ 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!
[newuser@localhost newuser]$
```

Figure 15–19 Removing and Restoring Permissions

A terminal window titled "Terminal" with a menu bar containing "File", "Edit", "Settings", and "Help". The terminal shows a user named "newuser" at "localhost" performing several commands. First, they run "chmod a-rw sneakers.txt", then "cat sneakers.txt", which results in a "Permission denied" error. Next, they run "chmod u+rw sneakers.txt", and then "cat sneakers.txt" again, which successfully displays the contents of the file: "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", and "relax!". The terminal ends with a prompt for the user.

```
[newuser@localhost newuser]$ chmod a-rw sneakers.txt
[newuser@localhost newuser]$ cat sneakers.txt
cat: sneakers.txt: Permission denied
[newuser@localhost newuser]$ chmod u+rw sneakers.txt
[newuser@localhost newuser]$ 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!
[newuser@localhost newuser]$
```

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`:

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

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`.

15.13 Fun with Numbers and `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 newuser newuser  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 access, but could still read the file (as shown in Figure 15-20, *Removing Group Write Permissions*), we'll have to remove the access by subtracting 2 from that set of numbers.

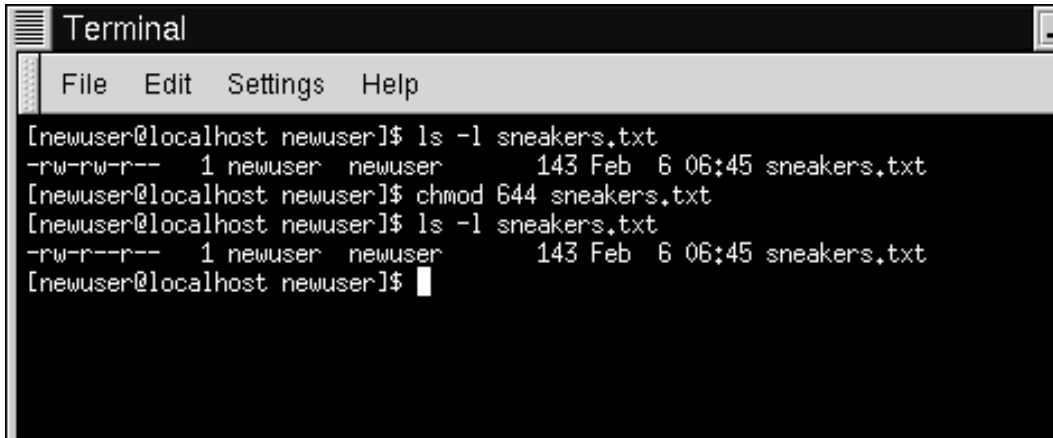
The numerical values, then, would become six, four and four -- or `644`.

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 newuser newuser  150 Mar 19 08:08 sneakers.txt
```

Figure 15–20 Removing Group Write Permissions

```
Terminal
File Edit Settings Help
[newuser@localhost newuser]$ ls -l sneakers.txt
-rw-rw-r-- 1 newuser newuser 143 Feb 6 06:45 sneakers.txt
[newuser@localhost newuser]$ chmod 644 sneakers.txt
[newuser@localhost newuser]$ ls -l sneakers.txt
-rw-r--r-- 1 newuser newuser 143 Feb 6 06:45 sneakers.txt
[newuser@localhost newuser]$
```

And there it is; now, neither the group nor others have write permission to `sneakers.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
```

WARNING

Biblical implications aside, setting permissions to 666 or 777 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.

16 Working with 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...

16.1 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 15, *You Are Here*, when we talked about redirection and piping, we were also demonstrating the power of `bash`.

More about bash

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 **newuser** might look something like this:

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

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

A standard for the PATH

Remember the earlier reference to the Filesystem Hierarchy Standard (see Section 15.4, *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.

16.2 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`.

More about `locate`

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.

More about `cron`

`cron` is a **daemon**. Daemons handle tasks in the background. 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.

16.3 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.

Looking at the environment variable

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].

Locating a previously used command

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 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).

16.4 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" is that file's extension).

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

16.4.1 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

16.4.2 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

16.4.3 System Files

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

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

16.4.4 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 15, *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`.

Read the man page

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 15, *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 15, *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.

Remember case

Unless otherwise specified, `grep` searches are **case sensitive**. 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 command.

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 example, just type:

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

This command behaves similarly to the command `ls -al /etc | more`. You may have used that command in Chapter 15, *You Are Here* to list the contents of the `/etc` directory then send the results through the `more` command for viewing on the screen.

It's safest to use >>

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.

Find out more about wildcards and regular expressions

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.

16.5 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.

16.5.1 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.
- `-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.

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.

```
[newuser@localhost newuser]$ 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 is the time to press [N] and [Enter].

16.5.2 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/newuser /home/newuser/tigger` using absolute pathnames.

16.5.3 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
```

16.5.4 Deleting Files and Directories

We talked about creating files with the `touch` command and by using redirection in Chapter 15, *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.

```
[newuser@localhost newuser]$ 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
```



The `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.

Part IV Q & A

17 Q & A: Quick Answers to Common Questions

17.1 Your First Login

17.1.1 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?

17.1.2 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**.

Forget your password?

If you forget your root password, don't panic. Turn to Section 17.11, *Unable to Log In* for instructions on how to repair the problem.

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 day-to-day tasks.) Refer to Section 1.1.1, *Create a User Account* for instructions on how to create your user account.

17.2 Using a Diskette

17.2.1 Q: Working with Diskettes

How do I use a diskette with Red Hat Linux?

17.2.2 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...

17.2.3 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.

17.2.4 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.

If you're using KDE

Want another quick way to format a floppy for either ext2 or MS-DOS? If you have KDE, try KDE Floppy Formatter, an easy way to format diskettes. To start the utility, go to **Main Menu K => Utilities => KFloppy**.

17.3 Switching Environments

17.3.1 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?

17.3.2 A: Using `switchdesk`

You can use a utility called **Switchdesk**, which allows you to change to and from GNOME, KDE and other environments.

Figure 17–1 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 17–1, *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.

17.4 Shutting Down

17.4.1 Q: "Right way" to Shut Down

What's the "right way" to shut down or reboot my system? Are there any shortcuts?

17.4.2 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.

17.5 Error Message While Installing an RPM

17.5.1 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.

17.5.2 A: Having "Permission"

If you're getting the error message similar to "failed to open /var/lib/rpm/packages.rpm" (see Figure 17-2, *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 17–2 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* on the Documentation CD (or online, at <http://www.redhat.com/support/manuals>).

17.6 Starting Applications

17.6.1 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?

17.6.2 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 executable 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.



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.

17.7 Accessing a Windows Partition

17.7.1 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?

17.7.2 A: Two Ways to Access Windows Partition.

You can gain access to another partition on your system -- for example, a Windows partition -- in a couple ways.

First, let's assume that your Windows partition is on your first IDE hard drive, in the first partition (`/dev/hda1`).

Open an Xterm window. If you're in your user account, `su` to root by typing

```
su
Password: yourrootpassword
```

Now, create a **mount point** to hold the data of your Windows partition, by typing

```
mkdir /mnt/vfat
```

To access the partition, as root in an Xterm, type the following:

```
mount -t vfat /dev/hda1 /mnt/vfat
```

Another method of mounting a Windows partition is by entering the correct fields in the file `/etc/fstab`. 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, following the above example.

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 not to mount the partition when you boot your system (**Not mount at boot time**), 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 access 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"`.

17.8 Find Some Commands Fast

17.8.1 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?

17.8.2 A: Searching `.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 [Space]; 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 17.10, *Tips on Using Command History*.

17.9 Keep `ls` Output from Scrolling

17.9.1 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?

17.9.2 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 a 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
```

17.10 Tips on Using Command History

17.10.1 Q: Tips and Tricks for `history`

What are some other ways I can use command history?

17.10.2 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.

17.11 Unable to Log In

17.11.1 Q: Using `linux single`

Help! I can't remember my root password. How do I log in now?

17.11.2 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.

17.12 Starting "Services"

17.12.1 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?

17.12.2 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, log in as root and type `ntsysv` at the shell prompt.

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.

17.13 Using All Available RAM

17.13.1 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?

17.13.2 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 LIInux 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.

17.14 Configuring a Sound Card

17.14.1 Q: How do I configure my sound card?

I can't get any system sounds, and I can't hear my audio CDs. How do I set up my sound card?

17.14.2 A: Use `sndconfig`

A tool called `sndconfig` can help you configure your sound card.

To use `sndconfig`, you must log in as root.

Once you're root, open an Xterm and type `sndconfig` at the prompt. This tool will help you configure your card for your Red Hat Linux system. To navigate through the **Yes**, **No** or **Cancel** buttons on the screens, just use your [Tab] and [Enter] keys.

Is it supported?

While many sound cards are supported for Red Hat Linux, not all are completely compatible -- or compatible at all. If you're having trouble configuring your sound card, check the Hardware Compatibility List at Red Hat's website to see if your card is supported. You'll find the list at <http://www.redhat.com/hardware>.

With `sndconfig`, you can probe your system for any cards. If the utility detects a plug and play sound card, it will automatically try to configure the correct settings by playing sound samples. If you can hear the samples, just select **Ok** when you're instructed and you're all set.

If the probe doesn't find any cards, you'll be presented with a list from which you can pick your card. Use the up and down arrow keys to scroll through the list. If your card is listed, highlight it, then press [Enter] (or [Tab] to the **Ok** button and press [Enter]).

Your next task will be to select the correct **I/O port**, **IRQ** and **DMA** settings. These settings depend on the jumper settings of the sound card. You can find information about these settings in your sound card documentation. If you share your machine

with Windows, you can also find your sound card's settings in the **Device Manager** tab in the **System** section of your **Control Panel**.

Once you've selected the right settings for your card, you'll be presented with sound samples. If you hear the samples, select **OK**.

You're done. After your card is configured, you'll return to the shell prompt. Just type `exit` to close the window, then log out of root and log back into your user account.

If `sndconfig` doesn't work

When `sndconfig` doesn't work -- the sample doesn't play and you still don't have audio sounds -- there are alternatives, although they're not quite as simple as `sndconfig`.



The following instructions are not recommended for complete novices.

First, you will need to know the values for your sound card. If you share your machine with Windows, you can find your settings in the **Device Manager** tab in the **System** section of your **Control Panel**. Under Linux, you must also check the following files to make certain that you don't have any devices using the sound card's resources:

- `/proc/pci`
- `/proc/interrupts`
- `/proc/ioports`
- `/proc/dma`

If yours is not a plug and play card, you can manually edit your `/etc/conf.modules` file to include the sound card module that it should use, for example:

```
alias sound sb
alias midi opl3
options opl3 io=0x388
```

```
options sb io=0x220 irq=7 dma=0,1 mpu_io=0x300
```

If you do have a plug and play card, the configuration is more complex, and may not work for all cards.

If you have already run `sndconfig`, you should have the file `/etc/isapnp.conf`. When `sndconfig` was run, the `isapnp` tools created this template file. This file can be hand edited a text editor, such as `pico`.

To become a little more familiar with the breakdown of this file, take a look at the man page for it by typing:

```
man isapnp.conf
```

There is also a great reference at the `isapnptools` homepage at <http://www.roe-stock.demon.co.uk/isapnptools/>

Using `sndconfig`

If you haven't run `sndconfig`, issue the command (as root): `/sbin/pnpdump > /etc/isapnp.conf` to create the file `isapnp.conf`, which you can edit.

When you have created the file `/etc/isapnp.conf`, open it with `PICO` by typing `pico /etc/isapnp.conf` (note that you don't have to actually be in `/etc` to edit the file).

For a little help in the layout of this file, the man page is an excellent guide to keep at hand for this. Each device will have a section that starts with about 4 lines that gives information on what the device is:

```
# Card 1: (serial identifier 6d ff ff ff ff f0 00 8c 0e)
# Vendor Id CTL00f0, No Serial Number (-1), checksum 0x6D.
# Version 1.0, Vendor version 1.0
# ANSI string -->Creative ViBRA16X PnP<--
```

Look for the section containing the sound card device. The lines in this section that you need to pay attention to are the lines within parentheses. These are the lines that need to be adjusted for your card's resources.

You make adjustments by **commenting** or **uncommenting** the beginning of the line, indicating that you want to use the resource specified on that line. (When you uncomment, you're making that line visible for use in the file; when you comment, it's invisible, for all intents and purposes.) Uncommenting is performed by removing the hash mark (#); to comment, just replace the hash mark. Once you have finished making these changes, look for the line towards the end of the file that says:

```
# (ACT Y)
```

Remove the "#," save the changes, and exit from Pico.

After setting up the new configuration file, and saving it, you need to activate it. Just type:

```
isapnp isapnp.conf
```

This will reinitialize the card with your newly changed values. If the values that you select the first time don't work, you can edit the file until you find values that don't conflict. Remember to activate the changes each time that you edit this file.

17.15 Removing Red Hat Linux

17.15.1 Q: Removing Red Hat Linux

How do I remove Red Hat Linux from my computer?

17.15.2 A. Remove LILO

To remove Red Hat Linux, you'll need to remove the LILO information from your Master Boot Record (MBR).

There are several methods to removing LILO from the MBR. Inside of Linux, you can replace the MBR with an earlier, saved version of the MBR using the following command:

```
/sbin/lilo -u
```

In DOS, NT and Windows 95, you can use `fdisk` to create a new MBR with the "undocumented" flag `/mbr`. This will *only* rewrite the MBR to boot the primary DOS partition. The command should look like:

```
fdisk /mbr
```

Removing partitions

If you want to remove Linux from a hard drive, and you attempted to use the default DOS `fdisk`, you will experience the "Partitions exist but they don't exist" problem. The best way to remove non-DOS partitions is with a tool that understands partitions *other than* DOS.

You can perform this with your Red Hat Linux installation floppy by typing **linux expert** at the `boot:` prompt, such as:

```
boot:linux expert
```

Select the option to install, rather than upgrade and, when you get to the partitioning section, choose `fdisk`. In `fdisk`, type `p` to print out your partition numbers; then remove the Linux partitions with the `d` command. When you're satisfied with the changes you've made, you can quit with a `w`, and your changes will be saved to disk. If you've deleted too much, type `q` to quit without changes.

When your partitions have been removed, you can reboot with `[Ctrl]-[Alt]-[Del]` rather than continue with the installation.

17.16 Changing Login from Console to X at Startup

17.16.1 Q: Changing to Graphical Login

How do I change my login from the console to the graphical screen?

17.16.2 A. Edit `/etc/inittab`

If you find yourself logging into your system at the console, only to type the `startx` command to start the X Window System, changing your login from console to graphical will save you a step or two.

You must edit one file, `/etc/inittab`, by changing just one number in the **run-level** section. When you're finished, log out and, the next time you log in, you'll have a graphical screen.

Open an Xterm window. If you're in your user account, `su` to root by typing

```
su
Password: yourrootpassword
```

Now, type `pico /etc/inittab` to edit the file with Pico. The file `/etc/inittab` will open. Within the first screen, you'll see a section of the file which looks like this:

```
# Default runlevel. The runlevels used by RHS are:
# 0 - halt (Do NOT set initdefault to this)
# 1 - Single user mode
# 2 - Multiuser, without NFS (The same as 3, if you do not have networking)
# 3 - Full multiuser mode
# 4 - unused
# 5 - X11
# 6 - reboot (Do NOT set initdefault to this)
#
id:3:initdefault:
```

To change from a console to a graphical login, you should change the number in the line `id:3:initdefault:` from a 3 to a 5.

WARNING

Change *only* the number of the default runlevel from 3 to 5.

Your changed line should look like:

```
id:5:initdefault:
```

When you're satisfied with your change, save and exit the file by using the [Ctrl]-[X] keys. You'll see a message telling you that the file has been modified, and asking you to confirm your change. Type [Y] for yes.

When the file is closed, exit from root by typing `exit`, and close the Xterm window by typing `exit`.

That's it. Your next login will be from the graphical screen.

17.17 Configuring X

17.17.1 Q: Changing Screen Resolution

Everything on my desktop looks tiny. How do I change my screen resolution?

17.17.2 A. Use Xconfigurator

After you have installed Red Hat Linux, you may wish to change your **screen resolution**, so that objects will fit better on your desktop. You can adjust the resolution with Xconfigurator, a tool which allows you to modify your X Window System settings.

To use Xconfigurator, you must log in as root. At a shell prompt, type `Xconfigurator` to start the application. The opening screen will look like Figure 17-3, *The Opening Dialog of XConfigurator*.

Other ways to start the utility

You can also start Xconfigurator two other ways -- both of which give you access to additional configuration tools. You must log in as root, remember.

From your application menu, such as the main menu button in GNOME, go to **System => Text mode tool menu**; or, from a shell prompt, type `setup`. Both methods will produce a menu of items which you can configure, such as X, your sound card, mouse and other features.

Figure 17–3 The Opening Dialog of XConfigurator



To navigate and make selections, use your [Tab] key. When your selection is highlighted, press the [Enter] key.

Is your hardware supported?

Make certain that your monitor and graphics card are supported for Red Hat Linux. To check, visit the Hardware Compatibility List at <http://www.redhat.com/hardware>.

To continue, tab until the **Ok** button is highlighted and press [Enter]. Xconfigurator will then probe your system for your video card, and suggest the correct X server to match your system, as in Figure 17–4, *Results of Video Card Probe*. Click the **Ok** button to proceed.

Figure 17–4 Results of Video Card Probe

Next, pick your exact monitor make and model number, as in Figure 17–5, *The Monitor Setup Dialog*. (Hint: You can often find the make and model number on the front or back of your monitor.)

Figure 17–5 The Monitor Setup Dialog



You'll have quite a few monitor choices -- the hash mark (#) indicates your position in the list. Use your up-arrow and down-arrow keys to move through the list, or use the [Home], [End], [Page Up] and [Page Down] keys. You can also jump to the manufacturer's place on the list by typing in the first letters of the name (such as ViewSonic) on your keyboard.

Monitor not on the list?

If you don't find your monitor among the listed makes and models, you can specify your monitor's horizontal and vertical sync frequencies by choosing **Custom** in the list. Refer to your monitor's documentation for information about its correct frequency rates. *Be careful!* Don't select a frequency range which exceeds your monitor's capacity, because you could damage or destroy it.


Once you have selected your monitor's make and model, a dialog will inform you that your card will be probed for its correct display settings. Don't worry if your monitor blinks during the process; it's normal.

If you don't want to probe

It isn't necessary to probe your video card. If you choose the **Don't Probe** option, the next dialog will ask you how much memory is on your video card, then you will be able to select your chosen resolution (see Figure 17-7, *Select Video Modes*).

Once the probing has finished, you'll find a dialog similar to Figure 17-6, *Suggested Video Settings*. If you want to accept the default settings, select **Use Default**. If you'd prefer different settings, select the **Let Me Choose** button.

Figure 17-6 Suggested Video Settings



If you want to choose video resolutions different than the default, you will be presented with the **Select Video Modes** dialog, such as Figure 17–7, *Select Video Modes*.

Figure 17–7 Select Video Modes



You can choose modes ranging from 8-bit to 24-bit, depending on the power of your computer and how many colors you want displayed. In each of the mode columns, you'll find resolution choices listed; using your [Tab] key, position the cursor in the boxes next to your chosen resolution and select by pressing [Space].

Choosing more than one resolution?

You can select more than one screen resolution -- for example, if you prefer **800x600** and **1024x768**, you can adjust your resolution "on the fly." To do this, use the [Ctrl]-[Alt]-[+ (plus)] or [Ctrl]-[Alt]-[- (minus)] keys.

When you're finished making your selections, press **Ok**.

Next, your video settings will be tested. You will see a graphical screen, with a dialog asking whether your settings appear correct. Following this dialog, another will ask whether you want a graphical login screen; if you answer **Yes**, you will have a graphical screen when you log in to your system. A final dialog will inform you of the location of your video configuration file (`XF86Config`, which is in `/etc/X11`).

If everything seems fine, log out, then log in again to your new screen resolution.

Part V Appendixes

A Glossary

Alpha

A RISC (*Reduced Instruction Set Computer*) architecture developed by Digital Equipment Corporation.

ATAPI

An abbreviation for *AT Attachment Packet Interface*. 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 *Basic Input/Output System*. 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 *Complex Instruction Set Computer*. 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 *Complementary Metal Oxide Semiconductor* -- 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 *Enhanced Integrated Drive Electronics*, 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 "Oops."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 *Fully Qualified Domain Name*. 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 *File Transfer Protocol*. 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 *Group 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 data storage device. 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 and read data to and 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 *Integrated Drive Electronics*, 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."

IMAP

Internet Message Access Protocol. A protocol used by certain mail servers.

Intel

Company responsible for producing the microprocessors that most commonly appear in PC-compatible personal computers. These processors include the 80386, 80486, and the Pentium line.

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.

ISP

Internet Service Provider.

Kernel

The central core 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.

Name Server

In TCP/IP networking terms, a name server 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 *Network File System*, 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 *Pluggable Authentication Modules*. 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 *Personal Computer Memory Card International Association*. 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.

PGP

An abbreviation for *Pretty Good Privacy*. PGP is an electronic privacy program which helps you ensure privacy by letting you encrypt files and e-mail. The encryption technology employed by PGP is very strong. PGP was created by Phil Zimmermann, and depends on **public key cryptography** for its effectiveness. Public key cryptography is a procedure in which users exchange "keys" to send secure documents to each other. For more on PGP, go to <http://www.pgp.com>.

PLIP

An abbreviation for *Parallel Line Internet Protocol*. PLIP is a protocol that permits TCP/IP communication over a computer's parallel port using a specially-designed cable.

POP

Post Office Protocol. A protocol used by certain mail servers.

POSIX

A somewhat mangled abbreviation for *Portable Operating System Interface*. A set of standards that grew out of the UNIX operating system.

PPP

Point-to-Point Protocol. A protocol that permits a dialup connection to an Internet Service Provider.

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 *Random Access Memory*. 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.

RAM disk

A virtual drive which uses part of a computer's RAM to store data. A RAM disk provides quick access to information, but unlike written data, that data in a RAM disk is lost when the computer is turned off.

Reboot

To restart the boot process. See also "Boot."

Red Hat, Incorporated

A software company which produces and markets software, 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 *Reduced Instruction Set Computer*. 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 *Read Only Memory*. 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 *RPM Package Manager*. RPM enables the installation, upgrading and removal of packages (software programs that together form the Red Hat Linux operating system).

SCSI

An abbreviation for *Small Computer System Interface*, 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 *Serial Line Internet Protocol*. SLIP is a protocol that permits TCP/IP communication over a serial line (typically over a dial-up modem connection).

SMB

Short for *Server Message Block*, SMB is the communications protocol used by Windows-based operating systems to support sharing of resources across a network. Also known as Samba.

SMTP

Simple *Mail Transfer Protocol*. A widely used definition for how to transfer mail over 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 (*Reduced Instruction Set Computer*) 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 demand 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 *Transmission Control Protocol/Internet Protocol*, TCP/IP is the networking standard used on the Internet today.

Torvalds, Linus

See Linus Torvalds.

UID

Short for *User 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.

XFree86

A free implementation of the X Window System.

² Just kidding -- thank you Ken Thompson and Dennis Ritchie of Bell Telephone Laboratories for your inspired operating system design!

B A Brief History of Linux and Red Hat

B.1 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...

B.1.1 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 (*RPM Package Manager*) 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.

C DOS to Linux Cheat Sheet

Many Linux commands you type at a shell prompt are not that different from the commands you would type either in MS-DOS or in Windows (from the MS-DOS prompt). In fact, some commands are identical.

C.1 DOS and Linux Commands

What follows are some common commands used at the MS-DOS prompt in Windows 9x, and in Linux, as well as a basic example of how the command is used at the Linux shell prompt. Note that these commands usually have a number of options. To learn more about each command, read its associated man page (for example, type **man ls** at the shell prompt to read about the **ls** command).

Table C–1 Similar Commands

Command's Purpose	MS-DOS	Linux	Basic Linux Example
Copies files	copy	cp	<code>cp thisfile.txt /home/thisdirectory</code>
Moves files	move	mv	<code>mv thisfile.txt /home/thisdirectory</code>
Lists files	dir	ls	<code>ls</code>
Clears screen	cls	clear	<code>clear</code>
Closes prompt window	exit	exit	<code>exit</code>
Displays or sets date	date	date	<code>date</code>
Deletes files	del	rm	<code>rm thisfile.txt</code>
"Echoes" output on the screen	echo	echo	<code>echo this message</code>

Command's Purpose	MS-DOS	Linux	Basic Linux Example
Edits files with simple text editor	edit	pico ¹	<code>pico thisfile.txt</code>
Compares the contents of files	fc	diff	<code>diff file1 file2</code>
Finds a string of text in a file	find	grep	<code>grep this word or phrase thisfile.txt</code>
Formats a floppy	format a: (if floppy's in A:)	mke2fs (or mformat ²)	<code>/sbin/mke2fs /dev/fd0</code> (<code>/dev/fd0</code> is the Linux equivalent of A:)
Displays command help	<code>command</code> /?	man ³	<code>man command</code>
Creates a directory	mkdir	mkdir	<code>mkdir directory</code>
Screens through a file	more	less ⁴	<code>less thisfile.txt</code>
Renames a file	ren	mv	<code>mv thisfile.txt that-file.txt⁵</code>
Shows your location in the file system	chdir	pwd	<code>pwd</code>

¹ Pico is a simple text editor; other editors you can use in place of pico include emacs and vi.

² This formats a disk for the DOS filesystem.

³ Or you can use info for some commands.

⁴ You can also another pager, called more, to scroll through a file a screen at a time.

⁵ The mv command serves double-duty, because it can both move a file and, if you want to rename a file in the same directory, you "move" that file to the same directory with a new name, as in this example.

Command's Purpose	MS-DOS	Linux	Basic Linux Example
Changes directories with a specified path (absolute path)	<code>cd path-name</code>	<code>cd path-name</code>	<code>cd /directory/directory</code>
Changes directories with a relative path	<code>cd ..</code>	<code>cd ..</code>	<code>cd ..</code>
Displays the time	<code>time</code>	<code>date</code>	<code>date</code>
Shows amount of RAM and use	<code>mem</code>	<code>free</code>	<code>free</code>

C.2 Different Paths

You'll notice other slight differences between the two operating systems, as well. Here is one of the more common examples:

C.2.1 Pathnames

A pathname shows you where you are, or conveys to the system which directory you want to go to. Briefly, the difference between stating the pathname is `\` or `/` when you're typing in the pathname.

For example, to navigate to a directory from the MS-DOS prompt, you would type

```
cd \thisdirectory\thisdirectory
```

To move to a directory (using absolute pathnames) in Linux, you would type

```
cd /thisdirectory/thisdirectory
```

Index

666 and 777..... 338

A

accounts

 modifying with Linuxconf 237

Alpha computer 412

appending standard output..... 322

applet, adding to the **Panel** in

 GNOME.....38

applications

 starting from shell prompt 369

 starting, on KDE desktop47

B

.bashrc file 306

boot disk

 creating

 mkbootdisk..... 287

C

cat 316

cd

 (See commands, cd)

chmod 327

 numerical settings..... 336

clear..... 315

color

 listing files in 306

command history..... 375

 tips 375

commands

 cat, using 316

 cd..... 294

 change directories..... 294

 chmod 327

 chmod, numerical settings 336

 clear..... 315

 copy..... 307

 cp..... 307

 history 373

 list contents 302

 login 363

 ls 302

 keeping output from

 scrolling 374

 ls -a 303

 ls -al..... 305

 ls -al --color..... 306

 ls, common options with..... 310

 mkdir 302

 print working directory..... 293

 pwd 293

 reset..... 315

 rm

 (See Files, deleting)

 rm -r

 (See directories, deleting)

 RPM, using 260

 stringing together 327

 su..... 301

 touch..... 302

common user questions 363

Communicator, configuring and

 using 129

Communicator, setting
 preferences..... 134
 Configuration
 Gnome-RPM 251

D

daemon, definition of.....75
 Desktop Guide, GNOME42
 desktop, in KDE49
 directories
 changing..... 294
 copying, moving, renaming 355
 deleting 357
 /etc 313
 listing contents 302
 skel..... 313
 /usr 313
 directories, managing at shell
 prompt..... 341
 diskettes
 creating an ext2 filesystem 365
 using 365
 DNS, definition of92
 documentation
 man pages at shell prompt 289
 dot files
 (See hidden files)
 drag and drop
 in GNOME33
 in KDE48
 drive icons in KDE50

E

e-mail

attaching documents 151
 attaching with "drag and drop". 152
 configuring your account in
 Messenger..... 138
 creating a signature 140
 reading and composing in
 Messenger..... 147
 return receipts in Messenger.. 154
 e-mail with Netscape
 Messenger..... 147
 Encrypting, decrypting with
 GnuPG..... 273
 environments
 changing..... 366
 changing, in GNOME.....43
 changing, in KDE.....63
 ext2 filesystem 365

F

FHS 311
 file manager
 GNOME, definition of36
 file managers
 GMC..... 217
 KFM 228
 file managers, GNOME and KDE 217
 files
 copying, moving, renaming 355
 deleting 357
 moving 356
 renaming 357
 files, managing at shell prompt ... 341
 filesystem
 understanding..... 311

- Filesystem Hierarchy Standard ... 311
 - finding commands 373
 - floppy diskette drive in KDE 50
 - floppy diskettes, using 365
 - folders, in KDE 52
- G**
-
- GIMP** 169
 - Bezier tool 172
 - brush
 - active 179
 - brushes 180
 - clone tool 190
 - color correction 188
 - color levels 188
 - colors 178
 - .gif images 183
 - gradient
 - active 179
 - image formats 183
 - image window 185
 - .jpg images 183
 - keyboard shortcut commands .. 187
 - move tool 174
 - opening a file 181
 - pattern
 - active 179
 - removing unwanted parts of an
 - image 190
 - saving a file 183
 - Script-Fu 210
 - selecting areas in an image 180
 - tiling a background 204
 - Toolbox 170
 - .xcf images 184
 - GMC** 36
 - using 217
 - GNOME**
 - desktop 33
 - folders in 35
 - help browser 40
 - latest documentation 31
 - launchers in 35
 - Pager 37
 - Panel 37
 - Panel, customizing 39
 - switching tasks 39
 - using mouse in 32
 - website location 31
 - workspace 33
 - GNOME and KDE, using application
 - from both 31
 - GNOME Control Center**
 - customizing desktop with 161
 - starting 161
 - GNOME in 60 Seconds 31
 - Gnome-RPM** 314
 - basic configuration 251
 - installing packages 249
 - main interface 247
 - package display 248
 - package manipulation 256
 - querying packages 256
 - selecting packages 248
 - uninstalling packages with 258
 - upgrading packages with 260
 - using 245–246
 - verifying packages 258

- GnoRPM
 - (See Gnome-RPM)
 - GNU Midnight Commander 217
 - viewing folders with.....36
 - Gnu Privacy Guard, using 273
 - GnuPG..... 273
 - generating a keypair 268
 - keypair, generating..... 268
 - graphical login screen, changing
 - to 384
- H**
-
- halt
 - safe ways to..... 367
 - halting and rebooting
 - from GNOME.....67
 - help
 - in KDE60
 - help browser in KDE60
 - help, finding, in GNOME.....40
 - help, searching in KDE62
 - hidden files 305
 - history
 - finding commands using..... 373
- I**
-
- image manipulation..... 169
- K**
-
- KDE and GNOME, using application
 - from both.....31
 - KDE Control Center
 - starting..... 165
 - KDE Control Center, customizing
 - desktop with 165
 - KDE Project45
 - KDE website location.....45
 - KDE, finding latest documentation.45
 - KDE, overview45
 - KFloppy..... 365
 - KFM 228
 - KLJetTool82
 - Kppp 108
- L**
-
- less 325
 - LILO
 - removing 383
 - Linuxconf
 - using 237
 - using to access Windows
 - partition 371
 - listing directories
 - (See commands, ls)
 - LJet Tool.....82
 - Log out button, adding, in GNOME.67
 - logging in 363
 - changing from console to
 - graphical 384
 - configuring.....19
 - logging out
 - at the console29
 - from GNOME.....65
 - from **Run program** in GNOME....68
 - in KDE68
 - killing an X session.....70
 - of root, shell prompt.....26

saving current settings.....69
 Login problems
 using single-user mode..... 376
 ls
 printing output..... 374
 viewing output 374
 Lynx 129

M

Main Menu Button in GNOME39
Main Menu K button58
 man pages
 at shell prompt..... 289
 man pages, reading with GNOME
 Help Browser.....40
 Messenger, configuring 138
 mkdir
 creating directories with 302
 mke2fs 365
 mounting a drive, definition51
 mouse
 using in KDE47
 mtools..... 364
 multiple desktops
 in GNOME41
 in KDE56

N

navigation 293
 Navigator
 keyboard shortcuts 133
 preferences..... 134
 using, in GNOME and KDE ... 130
 Netscape Communicator

 configuring..... 129
 setting preferences 134
 Netscape Messenger
 e-mail..... 147
 Networking, using RP3.....91
 news
 configuring your account in
 Messenger..... 138
 specifying servers in
 Messenger..... 145
 subscribing to newsgroups..... 158
 using **Messenger** 155
 ntsysv 377

O

online, connecting with RP3.....91
 ownership and permissions 327

P

package manipulation with
 Gnome-RPM 256
 packages
 dependencies 262
 freshening with RPM 264
 Gnome-RPM 249
 installing 261
 querying..... 264
 removing 262
 uninstalling with **Gnome-RPM** 258
 upgrading 263
 verifying..... 265
 verifying with **Gnome-RPM**... 256,
 258

- Pager
 - in GNOME37
 - paggers
 - less 325
 - Panel**
 - adding an application, in KDE...54
 - hiding, in KDE54
 - in GNOME37
 - in KDE54
 - in KDE, customizing54
 - parallel printer, configuring71
 - partitions
 - accessing a Windows 371
 - passwords, creating secure24
 - PATH**
 - editing 369
 - pathnames
 - relative and absolute 294
 - permissions
 - numerical settings..... 336
 - setting for new RPMs 368
 - permissions and ownership 327
 - pico
 - using to edit .bashrc 308
 - pipes 325
 - PPP91
 - configuring Kppp 110
 - configuring with **Dialup Configuration Tool**94
 - connecting with Kppp 108
 - connecting with RP3.....92
 - customizing RP3 102
 - Kppp, login script with 116
 - RH PPP Dialer**92
 - RP3, docking92
 - serial line names in MS-DOS and Linux 119
 - troubleshooting an RP3
 - configuration 106
 - pppd, definition of 102
 - printer
 - adding, in Printtool.....74
 - basic configuration71
 - compatibility71
 - /etc/printcap72
 - modifying configured82
 - monitoring jobs82
 - selecting the right filter78
 - The Linux Printing HOWTO.....71
 - using Printtool72
 - Printer icon in KDE**50
 - Printtool**72
 - pwd
 - (See commands, pwd)
- Q**
-
- Q & A 363
 - accessing a Windows partition
 - using Linuxconf 371
 - changing environments..... 366
 - configuring system services 377
 - finding previous commands 373
 - history tips and tricks..... 375
 - keeping `ls` output from
 - scrolling 374
 - login problems 376
 - permissions for installing
 - RPMs 368

recognizing RAM..... 377
 shutting down, rebooting 367
 sound card, configuring 380
 starting applications 369
 using large amounts of RAM... 377
 your first login..... 363
 querying packages with
 Gnome-RPM 256

R

Reboot
 safe ways to..... 367
 rebooting
 a shortcut..... 368
 Red Hat Linux
 history behind 411
 redirecting standard input..... 324
 redirection..... 318
 removing packages with
 Gnome-RPM 258
 removing Red Hat Linux 383
 reset..... 315
 root..... 311
 and root login..... 311
 using su..... 246
 RP3, connecting with.....92
 RPM..... 314
 and GnuPG 267
 checking package signatures ... 267
 dependency error..... 262
 errors on installing..... 263
 file conflicts, resolving..... 262
 freshen..... 264
 freshening packages 264

 from the shell prompt 260
 generating a keypair with
 GnuPG..... 268
 importing Red Hat's public key 270
 installing 261
 querying..... 264
 uninstalling 262
 upgrading 263
 using Gnome-RPM..... 246
 verifying..... 265
 RPM Package Manager 245, 314
 using from shell prompt..... 260
 RPM, using..... 245
 RPMs
 error message while installing . 368
 installing with Gnome-RPM ... 368

S

Screen resolution, changing 386
 selecting multiple items in KDE....47
 services
 starting of 377
 shut down
 safe ways to..... 367
 shutdown command
 in GNOME69
 shutdown command, options70
 single-user mode..... 376
 sndconfig 380
 SPARC computer 412
 standard input 316
 redirecting..... 324
 standard output 316
 appending 322

redirecting..... 318
 startup messages
 dmesg | more..... 326
 startx.....21
 superuser
 (See commands, su)
 Switchdesk
 using 366
 Switchdesk, in GNOME.....43
 switching desktops in KDE.....57
 switching tasks
 in KDE59
 system services
 configuring at startup..... 377

T

Taskbar, in KDE59
Tasklist, in GNOME43
Templates folder, in KDE.....52
 text editors
 Pico 308
 Torvalds, Linus 411
 touch
 creating files with..... 302
Trash icon in KDE.....50

U

upgrade 412
 Usenet, using Netscape
 Messenger..... 155
 user account, creating20
 user account, importance of
 creating.....19

utilities
 cat 316
 less 325

V

verifying packages with
 Gnome-RPM 258

W

window managers, definition of42
 Windows
 accessing on a separate partition
 adding line to /etc/fstab 371

X

Xterm, launchers in GNOME and
 KDE23