ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Red Hat Linux Bible: Fedora and Enterprise Edition

by Christopher Negus

ISBN:0764543334

John Wiley & Sons © 2003

Use this comprehensive guide to harness the power of Red Hat Linux on a server, desktop, or laptop computer. Coverage includes desktop and server basics, simple, painless installations, Linux Shell, using VmWare, VNC, and more.

Table of Contents

Red Hat Linux Bible—Fedora and Enterprise Edition
Preface
Part I - Getting Started in Red Hat Linux
Chapter 1 - An Overview of Red Hat Linux
Chapter 2 - Installing Red Hat Linux
Part II - Using Red Hat Linux
Chapter 3 - Getting Started with the Desktop
Chapter 4 - Using Linux Commands
Chapter 5 - Accessing and Running Applications
Chapter 6 - Publishing with Red Hat Linux
Chapter 7 - Playing Games with Red Hat Linux
Chapter 8 - Multimedia in Red Hat Linux
Chapter 9 - Tools for Using the Internet and the Web
Part III - Administering Red Hat Linux
Chapter 10 - Understanding System Administration
Chapter 11 - Setting Up and Supporting Users
Chapter 12 - Automating System Tasks
Chapter 13 - Backing Up and Restoring Files
Chapter 14 - Computer Security Issues
Part IV - Red Hat Linux Network and Server Setup
Chapter 15 - Setting Up a Local Area Network
Chapter 16 - Connecting to the Internet
Chapter 17 - Setting Up a Print Server
Chapter 18 - Setting Up a File Server
Chapter 19 - Setting Up a Mail Server
Chapter 20 - Setting Up an FTP Server
Chapter 21 - Setting Up a Web Server
Chapter 22 - Setting Up a News Server
Chapter 23 - Setting Up Boot Servers: DHCP and NIS
Chapter 24 - Setting Up a MySQL Database Server
Chapter 25 - Making Servers Public with DNS
Chapter 26 - Using Linux Servers from a Mac
Appendix A - What's on the CD-ROMs
Appendix B - Red Hat Linux RPMs
Appendix C - Running Network Services
Index
List of Figures

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

List of Tables
List of Sidebars

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Back Cover

If Red Hat Linux can do it, you can do it too…

With Red Hat Linux separating into the open source community Fedora Project and the commercial Red Hat Enterprise Linux 3, new opportunities arise for Red Hat Linux users. *Red Hat Linux Bible* leads you through the possibilities and choices between Fedora and enterprise installations. Covering everything from detailed instructions to running desktop applications and setting up more than a dozen server types, this book also features new sections on shell scripting, encryption techniques, setting up RAID disks, using yum, and much more.

Inside, you'll find complete coverage of Red Hat Linux

- Install, tune, and configure Fedora and Red Hat Linux Enterprise 3
- Navigate GNOME and KDE desktops to run the latest applications
- Learn to use the Linux shell, file system, and text editors
- Try out the latest security techniques for detecting and dealing with attacks and setting up encryption keys
- Discover how to install extra software packages to play games, enhance security, and administer Linux
- Install Linux on a laptop and manage power events with acpid
- Use the newest Linux multimedia tools for audio, video, and CD burning
- Configure Samba file/printer sharing, iptables, firewalls, CUPS printers, and other features using Red Hat as an AppleTalk or NFS server for Mac OS 9 or OS X computers
- Manipulate file systems, tune disks, and support multiple users

About the Author

Christopher Negus has been working with UNIX systems, the Internet, and (more recently) Linux systems for more than two decades. During that time, Chris worked at AT&T Bell Laboratories, UNIX System Laboratories, and Novell, helping to develop the UNIX operating system. Features from many of the UNIX projects Chris worked on at AT&T have found their way into Red Hat and other Linux systems.

Most recently, Chris co-authored the book *Linux Toys* for Wiley Publishing. During the past few years, Chris has written several books on UNIX and the Internet, including *Caldera OpenLinux bible, Internet Explorer 4 Bible,* and *Netscape Plug-Ins For Dummies* for Wiley Publishing. He also wrote several books for Que Corporation, including *The Complete Idiot's Guide to Networking* (second and third editions) and *Using UNIX* (second edition). Chris's other writings include articles for *Internet World, NetWare Connection*, and *Visual Developer* magazines.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Red Hat Linux Bible—Fedora and Enterprise Edition

Christopher Negus

Published by
Wiley Publishing, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright 2004 by Wiley Publishing, Inc. All rights reserved.

Published simultaneously in Canada
No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4744. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4447, E-Mail: permcoordinator@wiley.com.
Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or to obtain technical support, please contact our Customer Care Department within the U.S. at 800-762-2974, outside the U.S. at 317-572-3993 or fax 317-572-4002.
Limited Warranty: (a)WPI warrants that the Software and Software Media are free from defects in materials and workmanship under normal use for a period of sixty (60) days from the date of purchase of this Book. If WPI receives notification within the warranty period of defects in materials or workmanship, WPI will replace the defective Software Media. (b) WPI AND THE AUTHOR OF THE BOOK DISCLAIM ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WITH RESPECT TO THE SOFTWARE, THE PROGRAMS, THE SOURCE CODE CONTAINED THEREIN, AND/OR THE TECHNIQUES DESCRIBED IN THIS BOOK. WPI DOES NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE SOFTWARE WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION OF THE SOFTWARE WILL BE ERROR FREE. (c) This limited warranty gives you specific legal rights, and you may have other rights that vary from jurisdiction to jurisdiction.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Wiley, the Wiley Publishing logo, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc., and/or its affiliates in the United States and other countries and may not be used without written permission. Red Hat and Fedora are trademarks of Red Hat, Inc. Linux is a registered trademark of Linus Torvalds. All other trademarks are the property of their respective owners. Wiley is not associated with any product or vendor mentioned in this book.
Library of Congress Control Number: ***2003112325***

ISBN: 0-7645-4333-4

10 9 8 7 6 5 4 3 2 1

1O/RR/RZ/QT/IN

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

About the Author

Christopher Negus has been working with UNIX systems, the Internet, and (more recently) Linux systems for more than two decades. During that time, Chris worked at AT&T Bell Laboratories, UNIX System Laboratories, and Novell, helping to develop the UNIX operating system. Features from many of the UNIX projects Chris worked on at AT&T have found their way into Red Hat and other Linux systems. Most recently, Chris co-authored the book *Linux Toys* for Wiley Publishing. During the past few years, Chris has written several books on UNIX and the Internet, including *Caldera OpenLinux Bible, Internet Explorer 4 Bible*, and *Netscape Plug-Ins for Dummies* for Wiley Publishing. He also co-wrote several books for Que Corporation, including *The Complete Idiot's Guide to Networking* (second and third editions) and *Using UNIX* (second edition). Chris's other writings include articles for *Internet World*, *NetWare Connection*, and *Visual Developer* magazines.

At home, Chris enjoys spending time with his wife, Sheree, and his boys, Caleb and Seth. His hobbies include soccer, singing, and exercising with Sheree.

Credits

Acquisitions Editor
Debra Williams Cauley
Development Editor
Sara Shlaer
Production Editor
Eric Newman
Technical Editor
Jason Luster
Copy Editor
C. M. Jones
Editorial Manager
Mary Beth Wakefield
Vice President and
Executive Group Publisher
Richard Swadley
Vice President and Executive Publisher
Bob Ipsen
Vice President and Publisher
Joseph B. Wikert
Executive Editorial Director
Mary Bednarek
Project Coordinator
Bill Ramsey
Indexing
Johnna VanHoose Dinse

As always, I dedicate this book to my wife, Sheree. This book would never have happened without her love and support.

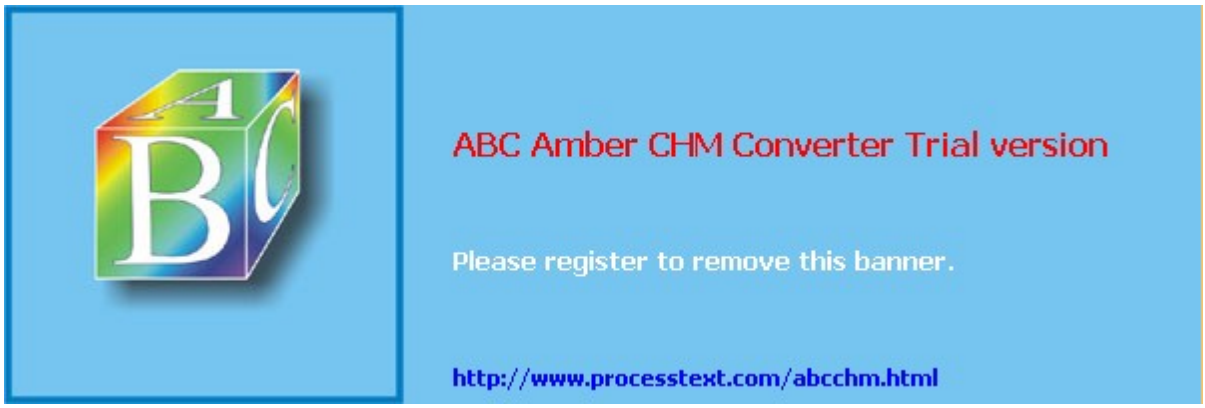

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Preface

Many Linux books resemble someone's throwing a bunch of high-performance car parts on the floor and saying, "Go ahead and build a Porsche." Although it's true that the parts you need for power computing are in Linux, you still need to know how to put them together. *Red Hat Linux Bible* takes you through those steps.

Who Are You?

You don't need to be a programmer to use this book. You may be someone who just wants to use Red Hat Linux (to run programs, access the Internet, and so on). Or you may simply want to know how to administer a Linux system in a workgroup or on a network.

I assume that you are somewhat computer-literate but have little or no experience with Linux (or UNIX). You may be migrating from Microsoft operating systems to Red Hat Linux because of its networking and multiuser features. You may be looking to start a career as a computer technician or network administrator and find that spending a few dollars for an entire operating system and book is more economical than taking those technical classes offered on late-night television. Or you might just think a "free" operating system is cool.

In any case, after you peruse this book you should have a good idea how to run applications, set up a small network, connect to the Internet, and configure a variety of server types (Web servers, print servers, file servers, and so on). This book represents a great first couple of steps towards your becoming someone who can set up a home network or a small office network and maintain a group of computers.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

This Book's Learn-through-Tasks Approach

The best way to learn a computer system is to get your hands on it. To help you learn Red Hat Linux, this book takes a task-oriented approach. Where possible, I step you through the process of working with a feature, such as setting up a network or configuring your desktop.

When you are done with a task, you should have a good, basic setup of the feature that it covers. After that, I often provide pointers to further information on tweaking and tuning the feature.
Instead of assuming that you already know about cryptic topics such as troff, NFS, and TCP/IP, I ease you into those features with headings such as "Publishing with Red Hat Linux," "Setting up a File Server," and "Connecting to the Internet." Heck, if you already knew what all those things were and how to get them working, you wouldn't need me, would you?

When many tools can be used to achieve the same results, I usually present one or two examples. In other words, I don't describe six different Web browsers, twelve different text editors, and three different news servers. I tell you how to get one or two similar tools really working and then note the others that are available.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

What You Need

To follow along with this book, you must install the software found on the accompanying CDs. To do that, you need a PC with the following general configuration:

An Intel Pentium or compatible CPU, 200Mhz or better (for text mode); 400Mhz or better (for GUI mode). (Some Intel 486 computers may work, but it's hit-or-miss.)

At least 64MB of RAM. To run the GNOME or KDE desktop 128MB are needed, although Red Hat recommends 256MB.

At least 510MB of hard disk space (you have to select a minimal install). You need 2.1GB of hard disk space for a personal desktop install, 2.6GB for a typical workstation installation, or at least 920MB of space for a server installation. To install everything, you need about 5.8GB of space.

A CD-ROM drive. This is recommended for installation, although you can install over a network or from a local hard disk instead. For those types of installs, you need at least a 3.5-inch floppy disk drive and either an extra hard disk partition or another computer (that can be reached over the network) that has packages or images of the Red Hat Linux CDs on it. (I tell you how to do that later, in case you're interested.)

Not every piece of PC hardware works with Red Hat Linux. You can find a list of computer processors that will work with Red Hat at `www.redhat.com/hardware`. You will probably want to use other types of hardware (also on that list) with your Red Hat Linux system, including video cards, mice, sound cards, modems, printers, scanners, joysticks, PCMCIA devices, and tape drives.

Cross-Reference I describe hardware requirements in more detail during descriptions of Red Hat Linux installation in Chapter 2.

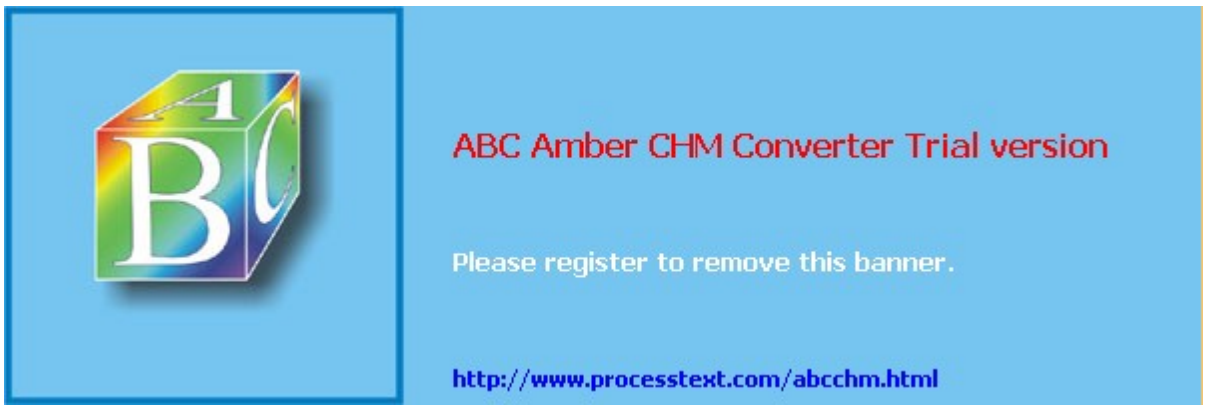

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Red Hat Linux Bible Improvements

Although this edition of *Red Hat Linux Bible* has *Fedora and Enterprise Edition* as the subtitle, you can think of it as reflecting a step further in the progression of Red Hat Linux. Until the name of the Linux distribution on the three CDs in this book was changed to Fedora Core, the distribution they reflected was just the next Red Hat Linux (Red Hat Linux 7.3, 8, 9, ... Fedora).
That said, going forward Red Hat, Inc. has made a more definitive split between the Fedora *Project* and Red Hat Linux Enterprise *Product*. For that reason, Chapter 1 is devoted to describing the differences between Fedora Core and Red Hat Enterprise Linux.

Because most of the technologies in Fedora and Enterprise have not drifted so far, most of the descriptions in this book will still work for both (with Enterprise including fewer of the bells, whistles, and toys you get with Fedora and more features that scale up to large servers.)

As this book, too is meant to reflect a progression of Red Hat Linux, I have included a variety of enhancements from the previous Red Hat Linux 9 Bible. Many of those enhancements are detailed below:

Installing Red Hat Linux — I adapted the installation procedure to match the Fedora Core CDs that come with this book, and noted where the installation process differs for Red Hat Enterprise Linux. I also noted that the LILO boot loader was removed from Fedora, although it is still in Enterprise for the time being.
Configuring RAID — I noted the ability to configure RAID disks during the installation process (Chapter 2). In Chapter 10, I go into more detail about how, and why, you would set up RAID disks in Red Hat Linux.
GNOME 2.4 — Although there are important changes for developers in GNOME 2.4, changes for end users are mostly reflected in a few new applications. In Chapter 9 I describe Epiphany, which replaces Galeon as the default Web browser for GNOME.
More text editors — For those who find the `vi` text editor challenging, in Chapter 4 I noted several other text editors you can try.
Using yum to get software — I added a description of the `yum` command and `yum.conf` file to Chapter 5. Yum is a great tool for downloading and installing add-on software packages in RPM format that are compiled specifically for each Red Hat Linux distribution (with dependent packages downloaded and installed as well).
Shift to ssh tools — Throughout this edition, I've tried to shift from old UNIX networking tools (`ftp`, `rlogin`, `rsh`, and the like) to more recent, more secure commands based on the OpenBSD secure shell protocol (`ssh`, `sftp`, and so on). Descriptions of `rlogin`, `rsh`, and other "r" commands have moved to the Wiley companion Web site for this book.
Using WineX 3.1 Point2Play — In Chapter 7, I added a description of the new Point2Play feature of WineX 3.1. Although not part of Red Hat Linux, this add-on can help you get Windows games running in Linux.
Cool Mozilla Web browsing tricks — Because we spend so much time browsing the Web these days, I like to add extra tips on using Mozilla to each edition. For this edition, I added some descriptions of tabbed browsing, the DOM inspector, and keystrokes for resizing Web pages to Chapter 9. Mozilla also has a new Junk Mail feature to help you build your own anti-spam rules in Mozilla Mail.
Improving laptop performance — A new section in Chapter 10 describes how to manage power better and generally improve laptop performance with Red Hat Linux.
Using up2date for software updates — I added some information on how to configure the new feature of up2date that lets you use apt and yum repositories to do software updates of Red Hat Linux software in the Fedora Core.
Removing temp files — I added a description to Chapter 11 about how tmpwatch automatically cleans up temporary files in Red Hat Linux.
Writing shell scripts — I pulled in my friend, and shell script expert, Kevin Pedigo to completely overhaul the section on writing shell scripts in Chapter 12. Along with new examples, Kevin added descriptions of how to use `grep`, `cut`, `tr` and `sed` in scripts.
Improving security — For Chapter 14, I had security expert Chuck Wolber add new sections on detecting and responding to denial-of-service, intrusion, and other attacks.
Creating certificates — Another major enhancement Chuck made to Chapter 14 was

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

a section on creating and obtaining SSL certificates to enable your Red Hat Linux system to provide secure communications.

A few new options — As I go through the book, I try to add new descriptions of options to old commands. A couple of note in this edition are the description of `rpm -F` in Chapter 5 (a convenient option to update many packages at once without installing new ones) and `umount -l` in Chapter 10 (to do lazy unmounts when resources are busy).

Using Linux NFS file sharing from the MAC — I added a procedure for configuring NFS in Linux (Chapter 26) so that a MAC OS X computer can use a shared NFS directory from the Linux machine.

Besides these sections, procedures throughout the book have been tested and corrected to match changes that have occurred to Red Hat Linux software in this version.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Conventions Used in This Book

Throughout the book, special typefaces indicate code or commands. Commands and code are shown in monospace type:

This is how code looks.

In the event that an example includes both input and output, monospace is still used, but input is presented in bold type to distinguish the two. Here is an example:

$ **ftp ftp.handsonhistory.com**

Name (home:jake): **jake**

Password: **********

The following boxes are used to call your attention to points that are particularly important.

Note A Note box provides extra information to which you need to pay special attention.

Tip A Tip box shows a special way of performing a particular task.

Caution A Caution box alerts you to take special care when executing a procedure, or damage to your computer hardware or software could result.

Cross-Reference A Cross-Reference box refers you to further information on a subject that you can find outside the current chapter.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

How This Book Is Organized

The book is organized into four parts.

Part I: Getting Started in Red Hat Linux

Part I consists of Chapters 1 and 2, which contain brief descriptions of the Linux technology and tell you what you need in order to get the operating system installed. Chapter 1 serves as an introduction to the Linux OS and to Red Hat Linux in particular. I also pay special attention to Red Hat, Inc.'s division between the Fedora Project and Red Hat Enterprise Linux. Chapter 2 discusses what you need in order to install Red Hat Linux (Fedora Core or Enterprise) and how to make the decisions you are faced with during installation. It includes procedures for installing from CD-ROM, hard disk, or network connection (NFS, FTP, or HTTP servers).

Part II: Using Red Hat Linux

Part II consists of Chapters 3 through 9, which include information for the average user who must use Linux to run applications and access the Internet. In Chapter 3, you learn about the GNOME desktop environment, the KDE desktop environment, and the X Window System. These GUIs provide graphical means of using Red Hat Linux. Chapter 4 describes ways of exploring and understanding Red Hat Linux, primarily from the Linux shell command interpreter. You learn how to use the `bash` shell, the vi text editor, and the commands for moving around the Linux file system.

Chapter 5 contains information on obtaining, installing, and running Linux applications. It also helps you run applications from other operating systems in Linux. Chapter 6 describes both old-time publishing tools and new, graphical word processors that are available with Red Hat Linux. Old tools include the troff and TeX text processing tools, whereas newer publishing software includes OpenOffice (included on the CDs) and StarOffice (commercially available).

Graphical and character-based games that run in Red Hat Linux are described in Chapter 7. This chapter also describes how to run commercial Windows games using WineX, and commercial Linux games, such as "Civilization: Call to Power" and "Myth II," some of which have demo versions available.

Chapter 8 describes how to use audio and video players, as well as how to configure sound cards and CD burners. Chapter 9 describes tools for browsing the Web (such as the Mozilla browser) and related tools (such as e-mail clients and newsreaders).

Part III: Administering Red Hat Linux

Part III consists of Chapters 10 through 14, which cover general setup and system maintenance tasks, including how to set up user accounts, automate system tasks, and back up your data. Chapter 10, in which you learn what you need to know about basic system administration, describes the root login, administrative commands, configuration files, and log files. Chapter 11 describes how to set up and provide support for multiple users on your Red Hat Linux system.

In Chapter 12 you learn to create shell scripts and to use the cron facility to automate a variety of tasks on your Red Hat Linux system. Techniques for backing up your system and restoring files from backup are described in Chapter 13. Chapter 14 describes issues related to securing your computing assets in Red Hat Linux.

Part IV: Red Hat Linux Network and Server Setup

Part IV consists of Chapters 15 through 26, which describe step-by-step procedures for setting up a variety of server types. Simple configurations for what might otherwise be complex tasks are contained in each chapter. Learn to arrange, address, and connect your Linux computers to a local area network (LAN) in Chapter 15. Chapter 16 describes techniques for connecting your Linux computer and LAN to the Internet, using features such as Point-to-Point Protocol (PPP), IP forwarding, IP masquerading, routing, virtual private networks, and proxy servers.

Chapter 17 describes how to set up different types of print server interfaces, including Samba (to share with Windows systems), NetWare, and native Linux printing. Chapter 18 describes file servers, such as Network File System (NFS) servers, Samba file servers, and NetWare file servers. Chapter 19 describes how to configure a sendmail e-mail server.

Chapter 20 describes how to configure and secure an FTP server, as well as how to access the server using FTP client programs. Chapter 21 teaches you how to set up Red Hat Linux as a Web server, focusing on the popular Apache server software. Chapter 22 describes how to set up an INN news server, including determining how to set up news feeds, choosing news storage methods, and enabling access to your INN server. Chapter 23 describes how to set up both DHCP and NIS services to distribute information to client workstations on the network.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 24 describes how to set up and use a MySQL database server in Linux. Chapter 25 takes you through the process of making the servers you configured in the other chapters available on the public Internet. Setting up a Domain Name System (DNS) server is also described in Chapter 25. Chapter 26 describes how to set up Red Hat Linux to be a Macintosh file and printer server.

Appendixes

This book contains three appendixes. Appendix A describes the contents of the companion CD-ROMs, Appendix B lists the hundreds of RPMs (software packages) that come with the Fedora Core distribution of Red Hat Linux, and Appendix C provides an overview of setting up and running network services.

Team LiB

PREVIOUS NEXT

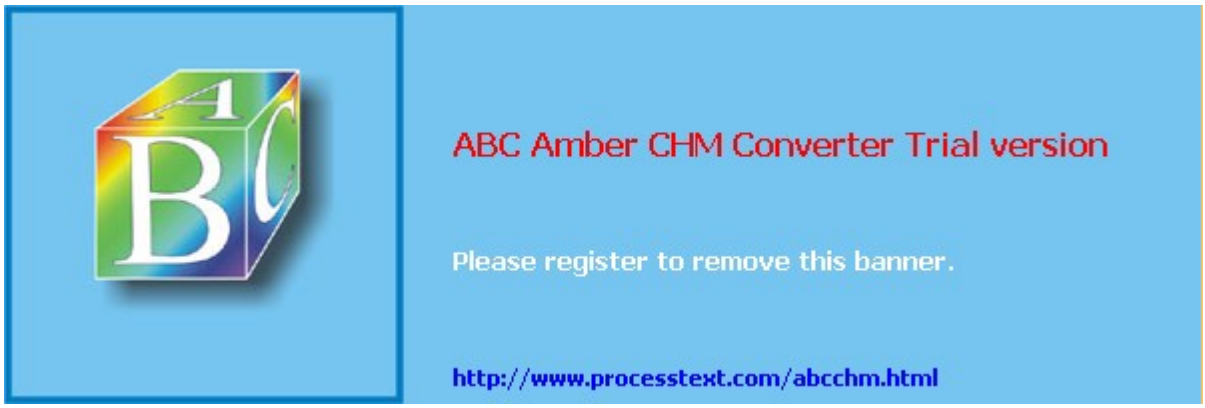

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

About the Companion CD-ROMs

The Fedora Core CD-ROMs that accompany this book provide the software you need for a complete, working Red Hat Linux system. With this software, you can install sets of software packages that result in personal desktop, workstation, or server configurations. This book describes how to configure and use the software for those different configurations.

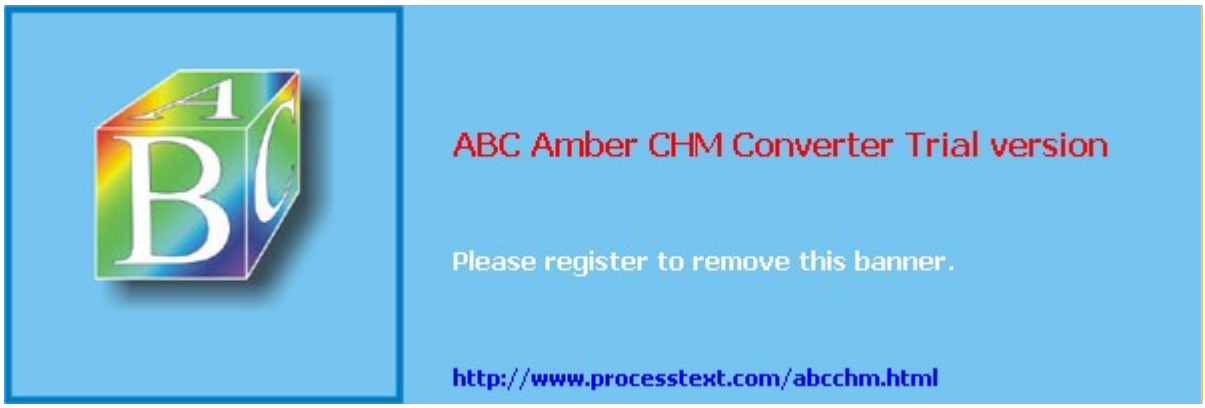

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

About the Companion Web Site

Even in a book that pushes the 1000-page boundary, there are a few topics that don't seem to make the cut. After you have gone through the book, you can visit the companion Web site at `www.wiley.com/compbooks/negus` for some bonus material on topics such as using the X Window system, finding neat add-on software, and exploring alternative administrative interfaces. Software moved to the Web site from the previous edition of this book includes information on using legacy UNIX remote commands (`rlogin`, `rcp`, and the like), running Tripwire to manage system security, and using crack and other password protection tools.

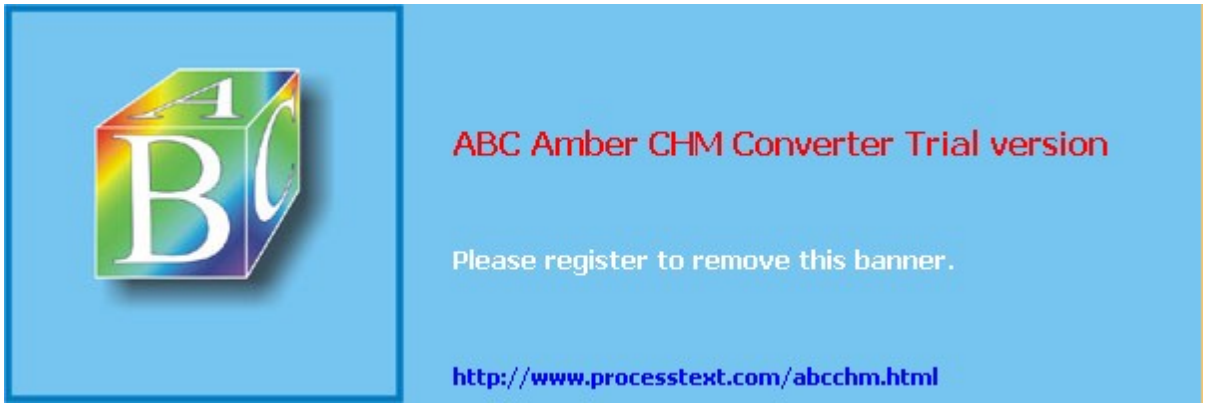

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Reach Out

If you have any questions or comments about this book, feel free to contact me by e-mail at this address: chris@linuxtoys.net.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Acknowledgments

A special acknowledgment to the people at Red Hat, Inc., who, despite the massive shift in direction of their Red Hat Linux product line, have managed to (once again) produce a rock-solid Linux distribution. We're proud to include it with this book. Great job!

At Wiley, I'd like to thank Debra Williams Cauley, who has continued her heroic efforts to push the latest editions of this book out to the world. Thanks to the tag team of Sara Shlaer and Eric Newman for their editing and production work on this book and on *Linux Toys* (so that I could get back to this book). Thanks to technical editor Jason Luster for plowing through the chapters yet again and to Margot Maley Hutchison and the others at Waterside Productions for bringing me this project.

For this edition, I'd like to acknowledge excellent contributions from:

Kevin Pedigo — For doing a great job rewriting the shell scripting section in Chapter 12.
Chuck Wolber — For adding new sections to Chapter 14 on identifying different kinds of security threats, as well as describing ways to combat those threats.
François Caen — For using his skills as a Red Hat systems administrator to help me sort through packaging differences in Fedora Core and Red Hat Enterprise Linux and understand the new yum and revised up2date tools.

I'd also like to thank members of the Tacoma Linux User's Group for being a great resource for discussions on Linux and (occasionally) non-Linux topics. On more than one occasion, the insights from its members have helped me fix or rewrite parts of this book.

Thanks, as always, to my dear family for helping me through this project. Hopefully, I'll be able to join Sheree, Caleb, and Seth on their next rock-climbing expedition.
Finally, a special thanks to those of you who bought this and earlier editions of *Red Hat Linux Bible*. Go out and become a force for Linux in your work, home and community. And if you're looking for something fun to do with Linux, check out *Linux Toys*, Wiley Publishing, Inc., by Christopher Negus and Chuck Wolber.

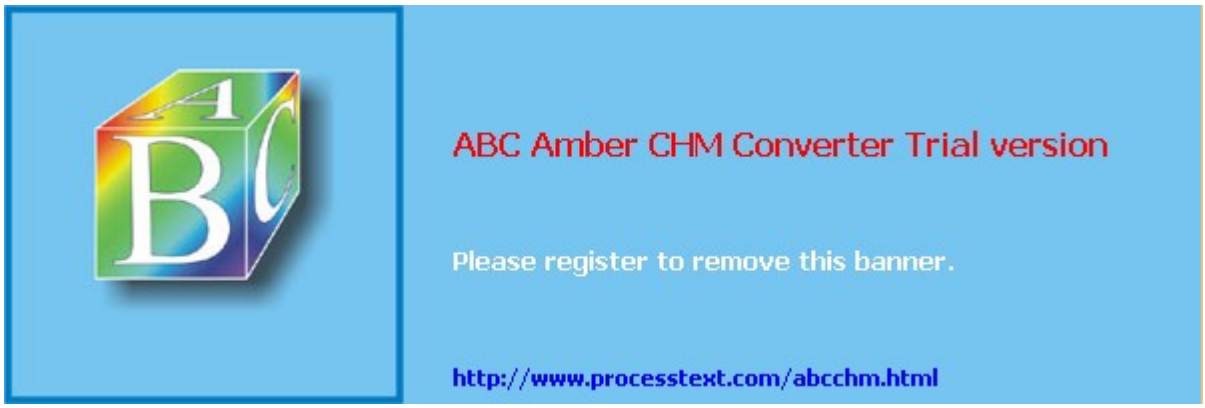

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Part I: Getting Started in Red Hat Linux

Chapter List

Chapter 1: An Overview of Red Hat Linux
Chapter 2: Installing Red Hat Linux

Team LiB

PREVIOUS NEXT

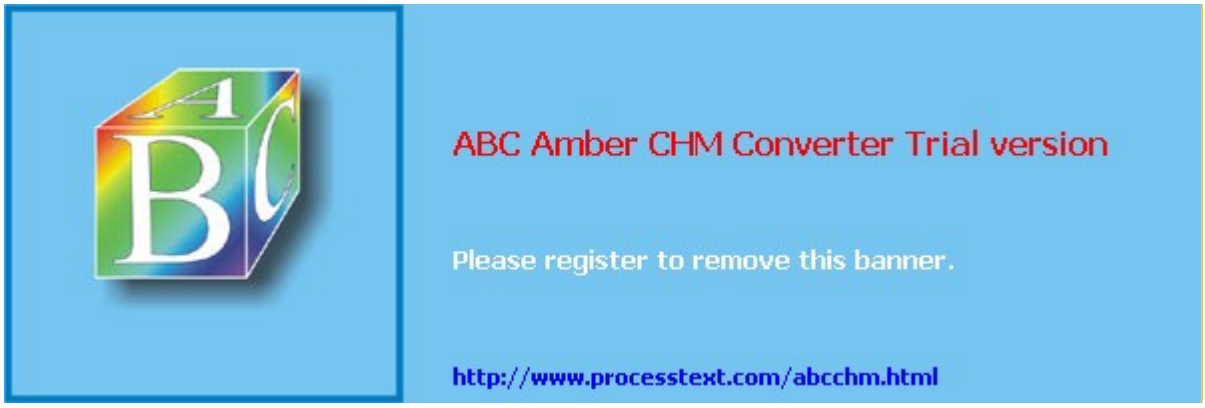

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 1: An Overview of Red Hat Linux

Overview

Linux was a phenomenon waiting to happen. The computer industry suffered from a rift. In the 1980s and 1990s, people had to choose between inexpensive, market-driven PC operating systems from Microsoft and expensive, technology-driven operating systems such as UNIX. Free software was being created all over the world, but lacked a common platform to rally around. Linux has become that common platform.

For several years, Red Hat Linux has been the most popular commercial distribution of Linux. With the latest versions of Red Hat Linux (reflected in the Fedora Core and Red Hat Enterprise Linux distributions), Red Hat, Inc. has taken steps to offer both free-flowing community versions and well-supported commercial versions of Red Hat Linux.

Note Because of significant overlap between Fedora Core and Red Hat Enterprise Linux, I use the term *Red Hat Linux* to refer to technology in both distributions. If software I describe is missing (primarily from Enterprise, which doesn't include many games and personal software), you can add the software later. Check your CDs, then check yum repositories described in Chapter 5 to find software RPMs.

Team LiB PREVIOUS NEXT

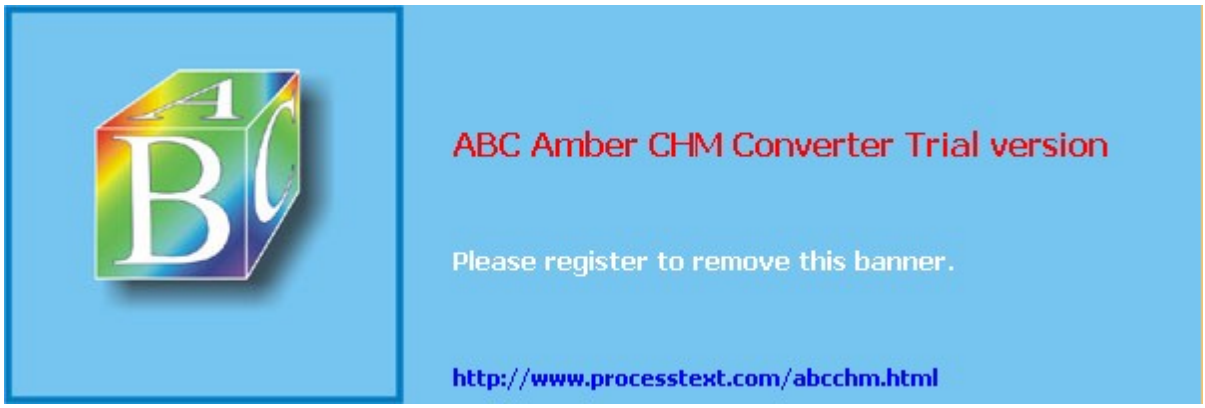

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Introducing Red Hat Linux

With the recent split between community (Fedora) and commercial (Red Hat Enterprise Linux) versions of Red Hat Linux, Red Hat has created a model that can suit the fast-paced changes in the open source world, while still meeting the demands for a well-supported commercial Linux distribution. (Later in this chapter I discuss differences between the Fedora and Enterprise versions of Red Hat Linux.)

Technical people have chosen Red Hat Linux because of its reputation for solid performance. With the new Fedora Project, Red Hat hopes to create an environment where open source developers can bring high-quality software packages to Red Hat Linux that would be beyond the resources of Red Hat, Inc. to test and maintain on its own.

Over 1,400 individual software packages (compared to just over 600 in Red Hat Linux 6.2) are included in the latest release of Red Hat Linux, referred to as Fedora Core. These packages contain features that would cost you hundreds or thousands of dollars to duplicate if you bought them as separate commercial products. These features let you:

- Connect your computers to a LAN or the Internet.
- Create documents and publish your work on paper or on the Web.
- Work with multimedia content to manipulate images, play music files, view video, and even burn your own CDs.
- Play games individually or over a network.
- Communicate over the Internet using a variety of Web tools for browsing, chatting, transferring files, participating in newsgroups, and sending and receiving e-mail.
- Protect your computing resources by having Red Hat Linux act as a firewall and/or a router to protect against intruders coming in through public networks.
- Configure a computer to act as a network server, such as a print server, Web server, file server, mail server, news server, and a database server.

This is just a partial list of what you can do with Red Hat Linux. Using this book as your guide, you will find that there are many more features built into Red Hat Linux as well.

Support for new video cards, printers, storage devices, and applications are being added every day. Linux programmers around the world are no longer the only ones creating hardware drivers. Every day more hardware vendors are creating their own drivers, so they can sell products to the growing Linux market. New applications are being created to cover everything from personal productivity tools to programs that access massive corporate databases.

Remember that old Pentium computer in your closet? Don't throw it away! Just because a new release of Red Hat Linux is out doesn't mean that you need all new hardware for it to run. Support for many old computer components get carried from one release to the next. There are old PCs running Red Hat Linux today as routers (to route data between your LAN and the Internet), firewalls (to protect your network from outside intrusion), and file servers (to store shared files on your LAN) — with maybe an Ethernet card or an extra hard disk added.

At this point, you may feel that Linux is something you want to try out. This brings us to the basic question: What is Linux?

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

What Is Linux?

Linux is a free operating system that was created by Linus Torvalds when he was a student at the University of Helsinki in 1991. Torvalds started Linux by writing a *kernel* — the heart of the operating system — partly from scratch and partly by using publicly available software. (For the definition of an operating system and a kernel, see the sidebar "What Is an Operating System?" later in this chapter.) Torvalds then released the system to his friends and to a community of "hackers" on the Internet and asked them to work with it, fix it, and enhance it. It took off.

Cross-Reference See Chapter 14 for a discussion about the difference between hackers (who just like to play with computers) and crackers (who break into computer systems and cause damage).

Today, there are hundreds of software developers around the world contributing software to the Linux effort. Because the source code for the software is freely available, anyone can work on it, change it, or enhance it. Developers are encouraged to feed their fixes and improvements back into the community so that Linux can continue to grow and improve.

On top of the Linux kernel effort, the creators of Linux also drew on a great deal of system software and applications that are now bundled with Linux from the GNU software effort (GNU stands for "GNU is Not UNIX"), which is directed by the Free Software Foundation (`www.gnu.org`). There is a vast amount of software that can be used with Linux, all of which includes features that can compete with or surpass those of any other operating system in the world.

If you have heard Linux described as a free version of UNIX, there is good reason for it. Although much of the code for Linux started from scratch, the blueprint for what the code would do was created to follow POSIX standards. POSIX (Portable Operating System Interface for UNIX) is a computer industry operating system standard that every major version of UNIX complied with. In other words, if your operating system was POSIX-compliant, it was UNIX. See the next section describing Linux's roots in the UNIX operating system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Linux's Roots in UNIX

Linux grew within a culture of free exchange of ideas and software. Like UNIX — the operating system on which Linux is based — the focus was on keeping communications open among software developers. Getting the code to work was the goal, without much concern about who owned the code, and the Internet was the primary communications medium. What, then, were the conditions that made the world ripe for a computer system such as Linux?

What Is an Operating System?

An operating system is made up of software instructions that lie between the computer hardware (disks, memory, ports, and so on) and the application programs (word processors, Web browsers, spreadsheets, and so on). At the center is the kernel, which provides the most basic computing functions (managing system memory, sharing the processor, opening and closing devices, and so on). Besides the kernel, an operating system provides other basic services needed to operate the computer, including:

File systems — The file system provides the structure in which information is stored on the computer. Information is stored in files, primarily on hard disks inside the computer. Files are organized within a hierarchy of directories. The Linux file system holds the data files that you save, the programs you run, and the configuration files that set up the system.
Device drivers — These provide the interfaces to each of the hardware devices connected to your computer. A device driver enables a program to write to a device without needing to know details about how each piece of hardware is implemented. The program opens a device, sends and receives data, and closes a device.
User interfaces — An operating system needs to provide a way for users to run programs and access the file system. Linux has both graphical and text-based user interfaces. GNOME and KDE provide graphical user interfaces, whereas shell command interpreters (such as `bash`) run programs by typing commands and options.
System services — An operating system provides system services, many of which can be started automatically when the computer boots. In Linux, system services can include processes that mount file systems, start your network, and run scheduled tasks. In Linux, many services run continuously, enabling users to access printers, Web pages, files, databases, and other computing assets over a network.

Without an operating system, an application program would have to know the details of each piece of hardware, instead of just being able to say, "open that device and write a file there."

In the 1980s while Microsoft flooded the world with personal computers running DOS and Windows operating systems, power users demanded more from an operating system. They ached for systems that could run on networks, support many users at once (multiuser), and run many programs at once (multitasking). DOS (Disk Operating System) and Windows didn't cut it.

UNIX, on the other hand, grew out of a culture where technology was king and marketing people were, well, hard to find. Bell Laboratories in Murray Hill, New Jersey, was a think tank where ideas came first and profits were somebody else's problem. A quote from Dennis Ritchie, co-creator of UNIX and designer of the C programming language, in a 1980 lecture on the evolution of UNIX, sums up the spirit that started UNIX. He was commenting on both his hopes and those of his colleagues for the UNIX project after a similar project called Multics had just failed:
What we wanted to preserve was not just a good environment in which to do programming, but a system

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

around which a fellowship could form. We knew from experience that the essence of communal computing as supplied by remote-access, time-shared machines, is not just to type programs into a terminal instead of a keypunch, but to encourage close communication.

In that spirit, the first source code of UNIX was distributed free to universities. Like Linux, the availability of UNIX source code made it possible for a diverse population of software developers to make their own enhancements to UNIX and share them with others.

By the early 1980s, UNIX development moved from the organization in Murray Hill to a more commercially oriented development laboratory in Summit, New Jersey (a few miles down the road). During that time, UNIX began to find commercial success as the computing system of choice for applications such as AT&T's telephone switching equipment, for supercomputer applications such as modeling weather patterns, and for controlling NASA space projects.

Major computer hardware vendors licensed the UNIX source code to run on their computers. To try to create an environment of fairness and community to its OEMs (original equipment manufacturers), AT&T began standardizing what these different ports of UNIX had to be able to do to still be called UNIX. To that end, compliance with POSIX standards and the AT&T UNIX System V Interface Definition (SVID) were specifications UNIX vendors could use to create compliant UNIX systems. Those same documents also served as road maps for the creation of Linux.
Today, Linux continues to aim toward POSIX compliance, as well as compliance with standards set by the new owner of the UNIX trademark, The Open Group (`http://www.unix-systems.org/`).

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Common Linux Features

No matter what version of Linux you use, the piece of code common to all is the Linux kernel. Although the kernel can be modified to include support for the features you want, every Linux kernel can offer the following features:

Multiuser — Not only can you have many user accounts available on a Linux system, you can also have multiple users logged in and working on the system at the same time. Users can have their own environments arranged the way they want: their own home directory for storing files and their own desktop interface (with icons, menus, and applications arranged to suit them). User accounts can be password-protected, so that users can control who has access to their applications and data.

Multitasking — In Linux, it is possible to have many programs running at the same time, which means that not only can you have many programs going at once, but that the Linux operating system can itself have programs running in the background. Many of these system processes make it possible for Linux to work as a server, with these background processes listening to the network for requests to log in to your system, view a Web page, print a document, or copy a file. These background processes are referred to as *daemons.*

Graphical User Interface (X Window System) — The powerful framework for working with graphical applications in Linux is referred to as the X Window System (or simply X). X handles the functions of opening X-based graphical user interface (GUI) applications and displaying them on an X server process (the process that manages your screen, mouse, and keyboard).

On top of X, you use an X-based desktop environment to provide a desktop metaphor and window manager to provide the look-and-feel of your GUI (icons, window frames, menus, and colors, or a combination of those items called *themes*). There are several desktop environments and several desktop managers to choose from. (Red Hat provides a few desktop managers, but focuses on GNOME and KDE desktop environments.)

Hardware support — You can configure support for almost every type of hardware that can be connected to a computer. There is support for floppy disk drives, CD-ROMs, removable disks (such as DVDs and Zip drives), sound cards, tape devices, video cards, and most anything else you can think of.

Note Not every hardware manufacturer provides Linux drivers with their peripheral devices and adapter cards. Although most popular hardware will be supported eventually in Linux, it can sometimes take a while for a member of the Linux community to write a driver.

Networking connectivity — To connect your Linux system to a network, Linux offers support for a variety of local area network (LAN) boards, modems, and serial devices. In addition to LAN protocols, such as Ethernet (both wired and wireless), all the most popular upper-level networking protocols can be built-in. The most popular of these protocols is TCP/IP (used to connect to the Internet). Other protocols, such as IPX (for Novell networks) and X.25 (a packet-switching network type that is popular in Europe), are also available.

Network servers — Providing networking services to the client computers on the LAN or to the entire Internet is what Linux does best. A variety of software packages are available that enable you to use Linux as a print server, file server, FTP server, mail server, Web server, news server, or workgroup (DHCP or NIS) server.

Application support — Because of compatibility with POSIX and several different application programming interfaces (APIs), a wide range of freeware and shareware software is available for Linux. Most GNU software from the Free Software Foundation will run in Linux (although some may take a bit of tweaking).

Note Because of the popularity of the Red Hat Package Management (RPM) format for packaging software, many software packages are available on the Internet in RPM format. If the RPM version matches your processor type (most have i386 and or i686 versions available), you can install the package without building and compiling the package. See Chapters 2 and 5 for information on working with RPM packages.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

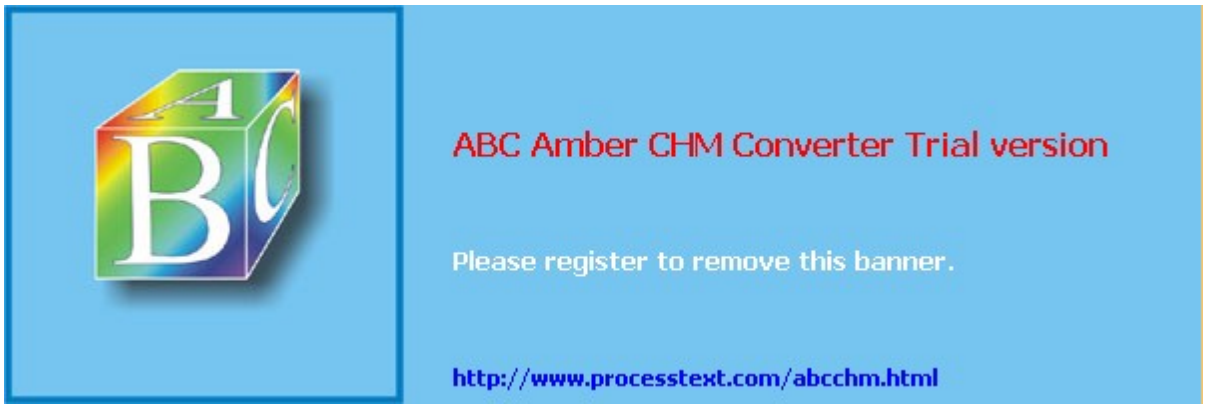

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Primary Advantages of Linux

When compared to different commercially available operating systems, Linux's best assets are its price and its reliability. Most people know that its initial price is free (or at least under $50 when it comes in a box or with a book). However, when people talk about Linux's affordability, they are usually thinking of its total cost, which includes the capability of using inexpensive hardware and compatible free add-on applications. Although commercial operating systems tend to encourage upgrading to later hardware, Linux doesn't (although faster hardware and larger disks are nice to have).

In terms of reliability, the general consensus is that Linux is comparable to many commercial UNIX systems but more reliable than most desktop-oriented operating systems. This is especially true if you rely on your computer system to stay up because it is a Web server or a file server. (You don't have to reboot every time you change something.)

Another advantage of using Linux is that help is always available on the Internet. There is probably someone out there in a Linux newsgroup or mailing list willing to help you get around your problem. Because the source code is available, if you need something fixed you can even patch the code yourself! On the other hand, I've seen commercial operating system vendors sit on reported problems for months without fixing them. Remember that the culture of Linux is one that thrives on people helping other people.

Note If you have general questions about Red Hat Linux, try the linux.redhat.misc newsgroup. For specific questions about networking or hardware, try the following newsgroups: comp.os.linux.networking and comp.os.linux.hardware, respectively.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

What Is Red Hat Linux?

Having directories of software packages floating extraneously around the Internet was not a bad way for hackers to share software. However, for Linux to be acceptable to a less technical population of computer users, it needed to be simple to install and use. Likewise, businesses that were thinking about committing their mission-critical applications to a computer system would want to know that this system had been carefully tested.
To those ends, several companies and organizations began gathering and packaging Linux software together into usable forms called *distributions.* The main goal of a Linux distribution is to make the hundreds of unrelated software packages that make up Linux work together as a cohesive whole. For the past few years, the most popular commercial distribution has been Red Hat Linux.

In September 2003, Red Hat, Inc., changed its way of doing business. That change resulted in the formation of the Red Hat–sponsored Fedora Project to take the development of Red Hat Linux technology into the future. But what does that mean to individuals and businesses that have come to rely on Red Hat Linux?

Red Hat forms the Fedora Project

The announcement of the Fedora Project by Red Hat, Inc. at first it prompted more questions than answers about the future direction of the company and its flagship Red Hat Linux product. In fact, it seemed that nothing named Red Hat Linux even existed anymore. Instead, what was Red Hat Linux would be reflected by Linux distributions coming from two paths:

Fedora Project (http://fedora.redhat.com) — An open source project, beginning from a Red Hat Linux 9 base, that produces its own Linux distribution. While the project is sponsored by Red Hat, Inc., there is no official support for the Linux distribution (called Fedora Core) that the project produces.
Red Hat Enterprise Linux (www.redhat.com/software/rhel) — An official, commercial Linux product from Red Hat, Inc. that is offered on an annual subscription basis. Red Hat backs up its Enterprise product line with technical support, training, and documentation.

The primary result of the Fedora Project (at least at first) is a set of three binary CDs and three source code CDs of a Linux distribution referred to as the Fedora Core. Before it was called Fedora Core, that distribution was being tested simply as the next in the series of Red Hat Linux distributions (presumably, Red Hat Linux 10). The three binary, installation CDs resulting from that effort are the same CDs that are included with this book.

The name change from Red Hat Linux to Fedora Core wasn't the only differences between the two distributions, however. Red Hat, Inc. also changed its association with the Fedora Core distribution in the following ways:

No boxed sets — Red Hat decided to not sell Fedora Core through retail channels. The ever-shortening release cycle was making it difficult to manage this, and Red Hat believed that most of its customers were clever enough to download the software.
No technical support offerings — There are no technical support programs available from Red Hat for Fedora Core, although at the moment you can still purchase the $60/year update service entitlement for Fedora.
No Red Hat documentation — The set of manuals that came with the previous Red Hat Linux product is not being brought over to Fedora. Instead, a series of small task-oriented documents will be collected for the project in article format.

By not creating a whole support industry around Fedora Core, that project is free to produce software release on a much shorter schedule (possibly two or three times per year). This allows Fedora Core users to always have the latest software features and fixes included with a recent version of the operating system.

Another potential upside to Fedora Core is that the Fedora Project hopes to encourage community software developers to create compatible software. By including software download and installation tools (such as the yum utility) in Fedora Core, the Fedora Project hopes to encourage people to contribute to

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

software repositories that Fedora Core users can rely on to download additional software packages.

Red Hat shifts to Enterprise Linux

The major shift of attention to Red Hat Enterprise Linux as the focus of Red Hat, Inc.'s commercial efforts has been on the horizon for some time. Some characteristics of Red Hat Enterprise Linux are:

Longer release intervals — Instead of offering releases every 4 to 6 months, Enterprise software will have a 12 to 18 month update cycle. Customers can be assured of a longer support cycle without having to upgrade to a later release.

Multiple support options — Customers will have the choice of purchasing different levels of support. All subscriptions will include the Update Module, which allows easy access to updates for Red Hat Enterprise Linux systems. The Management Module lets customers develop custom channels and automate management of multiple systems. The Monitoring Module allows customers to monitor and maintain an entire infrastructure of systems.

Documentation and training — Manuals and training courses will center on the Red Hat Enterprise Linux distribution.

Red Hat Enterprise Linux install types focus on three different types of computer systems, referred to as WS (for workstations), AS (for high-end systems), and ES (for small/mid-range servers). Each system in the Red Hat Enterprise Linux family is meant to be compatible with the others. There are Basic, Standard and Premium editions of these Enterprise systems. While Basic offers only software downloads, standard and premium editions offer hard copy documentation and additional technical support.

Choosing between Fedora and Enterprise

If you bought this book to try out Linux for the first time, rest assured that what you have on the three CDs with this book is a solid, battle-tested operating system. Because (at the moment at least) there is so much overlap between Fedora Core and Red Hat Enterprise Linux, Fedora Core provides a way to test out much of the software that is in Enterprise editions.

Although Fedora Core may not be right for everyone, Fedora Core is still great for students, home users, most small businesses, and anyone just wanting to try out Red Hat Linux technology. Larger businesses should seriously consider the implications to support, training, and future upgrade paths before choosing whether to go the Fedora route or sign on with Red Hat Enterprise Linux.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Why Choose Red Hat Linux?

To distinguish themselves from other versions of Linux, each distribution adds some extra features. Because many power features included in most Linux distributions come from established open source projects (such as Apache, Samba, KDE, and so on), often enhancements for a particular distribution exist to make it easier to install, configure, and use Linux. Also, because there are different software packages available to do the same jobs (such as window managers or a particular server type), a distribution can distinguish itself by which packages it chooses to include and feature.

Red Hat Linux traditionally has set itself apart from other Linux distributions with these features:

Software packaging — Red Hat, Inc. created the Red Hat Package Management (RPM) method of packaging Linux. RPMs allow less technically savvy users to easily install Linux software. With RPM tools, you can install from CD, hard disk, over your LAN, or over the Internet. It's easy to track which packages are installed or to look at the contents of a package. Because RPM is available to the Linux community, it has become one of the de facto standards for packaging Linux software.

Cross-Reference Chapter 5 describes how to install RPM packages.

Easy installation — The Red Hat Linux installation process (called anaconda) provides easy steps for installing Linux. During installation, Red Hat also helps you take the first few steps toward configuring Linux. You can choose which packages to install and how to partition your hard disk. You can even get your desktop GUI ready to go by configuring your video card, user accounts, and even your network.

Cross-Reference Chapter 2 covers Red Hat Linux installation.

UNIX System V–style run-level scripts — To have your system services (daemon processes) start up and shut down in an organized way, Red Hat Linux uses the UNIX System V mechanism for starting and stopping services. Shell scripts (that are easy to read and change) are contained in subdirectories of `/etc`. When the run level changes, such as when the system boots up or you change to single-user mode, messages tell you whether each service started correctly or failed to execute properly. Chapter 12 describes how to use run-level scripts.

Desktop environments (GNOME and KDE) — To make it easier to use Linux, Red Hat Linux comes packaged with GNOME and KDE desktop environments. GNOME is installed by default and offers some nice features that include drag-and-drop protocols and tools for configuring the desktop look and feel. KDE is another popular desktop manager that includes a wide range of tools tailored for the KDE environment, such as the KDE Control Center for configuring the desktop.

Desktop look-and-feel — With the latest Red Hat Linux distributions, whether you use KDE or GNOME as your desktop environment, you can expect to see many of the same icons and menus to help standardize how you use your Red Hat Linux system. Tools you can launch from those environments help you configure your network, set up servers, watch log files, and manage system services.

GUI Administration tools — There are some helpful Red Hat configuration tools for setting up some of the trickier tasks in Linux. Several different GUI tools provide a graphical, form-driven interface for configuring networking, users, file systems, and initialization services. Instead of creating obtuse command lines or having to create tricky configuration files, these graphical tools can set up those files automatically.

Note There are advantages and disadvantages of using a GUI-based program to manipulate text-based configuration files. GUI-based configuration tools can lead you through a setup procedure and error-check the information you enter. However, some features can't be accessed through the GUI, and if something goes wrong, it can be trickier to debug. With Linux, you have the command-line options available as well as the GUI administration tools.

Testing — The exact configuration that you get on a Red Hat Linux distribution has been thoroughly tested by experts around the world. The simple fact that a software package is included in the Red Hat Linux distribution is an indication that Red Hat

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

believes it has achieved a certain level of quality. Although Fedora Core will transition to a community-based approach to testing, at least for the first release Fedora Core was put through the same basic testing process as previous versions of Red Hat Linux.

Automatic updates — The software packages that make up Red Hat Linux are constantly being fixed in various ways. To provide a mechanism for the automatic selection, download, and installation of updated software packages, Red Hat created the Red Hat Network. Using the Red Hat Network Web site or the `up2date` command, you can receive critical security fixes and patches very simply over the Internet. Fedora Core still allows automatic updates through Red Hat Network for the time being.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Features in Red Hat Linux

The major components in the Fedora Core include (with version numbers):

Linux kernel: version 2.4.22
GNOME (desktop environment): version 2.4
KDE (desktop environment): version 3.1.4
GCC (GNU C language compilation system): version 3.3
Apache (Web server): version 2.0.47
Samba (Windows SMB file/printer sharing): version 3.0
CUPS (print services): version 1.1.19
Sendmail (mail transport agent): version 8.12
vsFTPd (secure FTP server): version 1.2
INN (Usenet news server): version 2.3.5
MySQL (database server): version 3.23.58
BIND (Domain name system server): version 9.2.2

Tip If you want the latest features in Linux when looking at different Linux distributions, compare the version numbers shown above. Version numbers and names that Linux distributors such as Mandrake, SUSE, and Red Hat associate with their releases can be arbitrary. By comparing versions of the kernel, KDE and GNOME desktops, and GNU compiler they are using, you can tell which distribution actually has the latest features.

As Red Hat continues to consolidate its distribution, some popular packages have been dropped from Fedora Core since the previous version of Red Hat Linux, such as the following:

LPRng — This printing service was dropped, making CUPS the preferred (and only) printing service included with Red Hat Linux.
exmh — This e-mail handling system was dropped from this release.
galeon — This Web browser was replaced by the epiphany Web browser.
pine — This e-mail reader was dropped due to maintenance and licencing concerns.
LILO — This boot loader has been replaced by GRUB by and removed from the distribution.
tripwire — This software for checking the integrity of your Red Hat Linux system was removed from this version. (You can find information on tripwire on the *Red Hat Linux Bible* Web site: www.wiley.com/compbooks/negus.)

Here are a few other packages that are noted as deprecated (slated to be dropped in the future):

Glide3 — This package supports 3D graphics for Mesa and 3Dfx Voodoo hardware.
sndconfig — Redhat-config-soundcard is recommended for configuring sound cards.
ncpfs— This package is used to set up NetWare servers in Linux.

See Appendix B for information on other packages no longer included in Red Hat Linux.

Note Just because a package has been dropped from Red Hat Linux doesn't mean that you can't still get and use the package. In fact, I tell how to find and install packages like wine and WuFTPd in this book that have been dropped from previous versions of Red Hat Linux.

The following paragraphs describe many of the major features in both Fedora Core and Red Hat Enterprise Linux.

Red Hat config tools

Red Hat has added to its growing arsenal of graphical administrative tools. Since dropping the linuxconf and bypassing the Webmin graphical administrative interfaces, Red Hat has been steadily developing and adding its own administrative tools to its distribution. As a result, a systems administration can often skip running shell commands and editing plain-text configuration files to set up servers, manage system resources, or add users.

The following is a list of Red Hat GUI packages and what each is used to configure:

redhat-config-bind: Domain Name System server.
redhat-config-date: System time and date.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

redhat-config-httpd: Apache Web server.
redhat-config-language: Languages for Red Hat Linux.
redhat-config-keyboard: Keyboard selection.
redhat-config-kickstart: Kickstart files for unattended Red Hat Linux installations.
redhat-config-mouse: A mouse.
redhat-config-network: Network interfaces.
redhat-config-nfs: Network File System shared directories.
redhat-config-packages: Red Hat Linux software.
redhat-config-printer: Printers.
redhat-config-printer-gui: Printers (GUI).
redhat-config-printer-tui: Printers (text-based).
redhat-config-proc: Kernel tunable parameters.
redhat-config-rootpassword: Change your root password.
redhat-config-samba: Samba Windows file/printer sharing (new in Red Hat Linux 9).
redhat-config-securitylevel: Iptables firewalls.
redhat-config-services: System services.
redhat-config-soundcard: Sound card.
redhat-config-time: Set time and date.
redhat-config-users: User accounts.
redhat-config-xfree86: X display and monitor.
redhat-logviewer: System log file viewer.

You can launch the tools associated with the previous packages either from the main Red Hat menu or from a Terminal window. In most cases, the name of the command you run to launch the window is the same name as the package it comes in.

Red Hat desktop interfaces

KDE and GNOME are desktop environments that provide a framework for running and developing graphical applications and offer a full range of preferences to allow users to tailor the exact desktop look-and-feel. There was some hubbub surrounding the new "Red Hat" look-and-feel implemented in Red Hat Linux 8, which gave both GNOME and KDE similar menus, icons, and colors. Die-hard GNOME and KDE users, however, are free to change their desktops however they like.

Unlike previous releases of Red Hat Linux, you have to work a bit to get the KDE desktop. If you choose to install Red Hat Linux as a Personal Desktop or Workstation system, you get the GNOME desktop by default. You must specifically ask to install additional packages to get KDE. By default, it is only included in an Everything install.

Additional software packages

By far, most of the enhancements to Fedora Core over previous versions of Red Hat Linux have come in existing packages. Some new packages have been added, however. Along with those mentioned in previous sections, the following list gives you an idea of some of the new software packages added to Red Hat Linux:

acpid — Provides tools for the Advanced Configuration and Power Interface feature.
Bluetooth — The bluez-bluefw, bluez-hcidump, bluez-pan, and bluez-sdp packages offer support for Bluetooth wireless networking.
brltty — Contains support for braille terminal access to console.
dovecot — Contains the dovecot secure IMAP server.
dvd+rw-tools — Contains tools for mastering DVD content.
epiphany — A Web browser that replaces the galeon Web browser in this release.
freeradius — Contains the remote authentication dial-in user service authentication protocol..
gnome-pilot-conduits — Contains the Palm OS PDA conduits.
nano — Contains the nano text editor.
yum — Contains the yum command for downloading and installing selected software packages in RPM format.

There are other new packages in this release of Fedora Core as well. For a complete list of packages in this release of Red Hat Linux (all of which are on the CDs that come with this book), see Appendix B.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

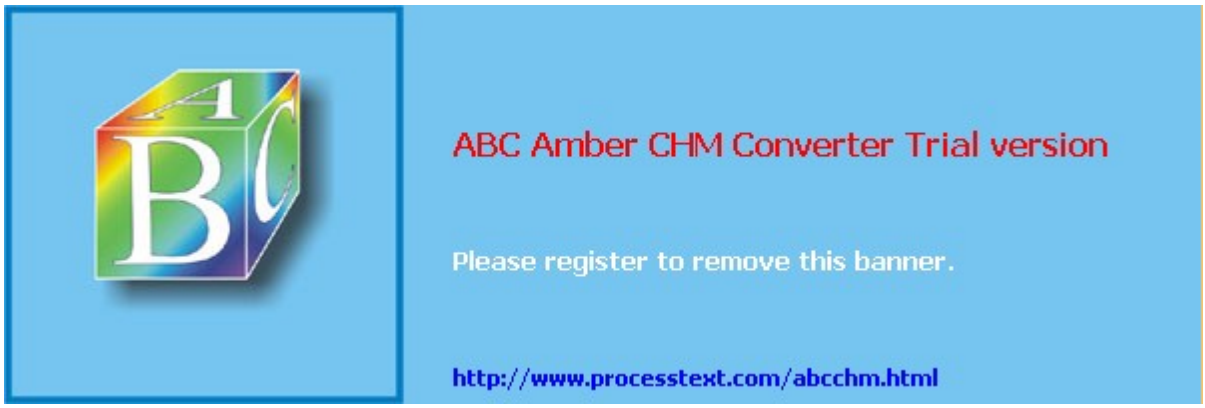

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The Culture of Free Software

I would be remiss to not say something about the culture of free software development from which Linux has thrived and will continue to thrive. The copyright for Red Hat Linux is covered under the GNU public license. That license, which most free software falls under, provides the following:

Author rights — The original author retains the rights to his or her software.

Free distribution — People can use the GNU software in their own software, changing and redistributing it as they please. They do, however, have to include the source code with their distribution (or make it easily available).

Copyright maintained — Even if you were to repackage and resell the software, the original GNU agreement must be maintained with the software. This means that all future recipients of the software must have the opportunity to change the source code, just as you did.

It is important to remember that there is no warranty on GNU software. If something goes wrong, the original developer of the software has no obligation to fix the problem. However, the Linux culture has provided resources for that event. Experts on the Internet can help you iron out your problems, or you can access one of the many Linux newsgroups to read how others have dealt with their problems and to post your own questions about how to fix yours. Chances are that someone will know what to do — maybe even going so far as to provide the software or configuration file you need.

Note The GNU project uses the term *free software* to describe the software that is covered by the GNU license. On occasion, you may see the term *open-source software* being used to describe software. Though source code availability is part of the GNU license, the GNU project claims that software defined as open source is not the same as free software because it can encompass semi-free programs and even some proprietary programs. See www.opensource.org for a description of open-source software.

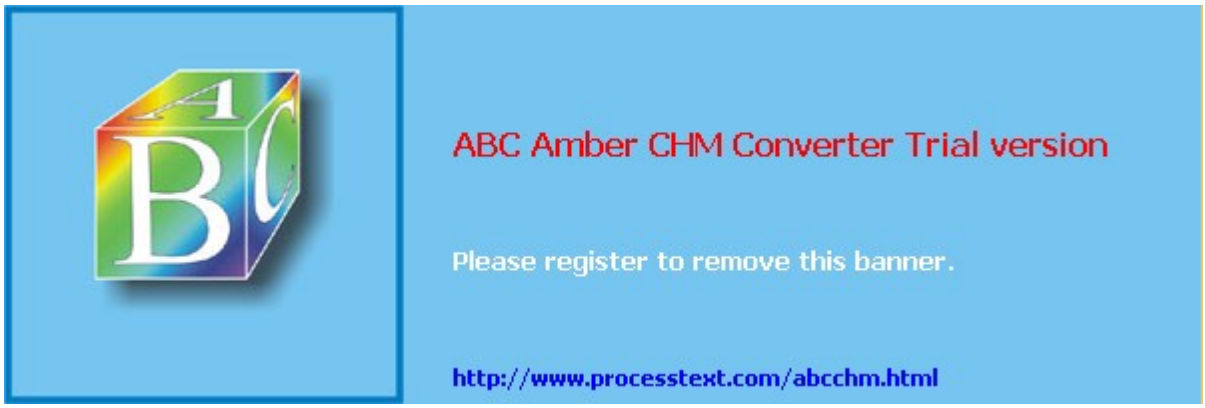

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

Linux is a free computer operating system that was created by Linus Torvalds in 1991 and has grown from contributions from software developers all over the world. Fedora Core and Red Hat Enterprise versions of Red Hat Linux are distributions of Linux that package together the software needed to run Linux and makes it easier to install and run.

Features in Red Hat Linux distributions include a simplified installation procedure, Red Hat Package Management (RPM) tools for managing the software, and easy-to-use GNOME and KDE desktop environments. You can get Red Hat Linux from the Internet or from distributions that come with books such as this one.

Linux is based on a culture of free exchange of software. Linux's roots are based in the UNIX operating system. UNIX provided most of the framework that was used to create Linux. That framework came in the form of the POSIX standard, which defined the interfaces a UNIX system needs to support in order to be a UNIX system.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 2: Installing Red Hat Linux

Overview

A simplified installation procedure is one of the best reasons for using a Linux distribution such as Red Hat Linux. In many cases, for a computer dedicated to using Red Hat Linux, you can just pop in the CDs (that come with this book), choose from several preset configurations, and be up and running Linux in less than an hour.

If you want to share your computer with both Linux and Microsoft Windows, Red Hat offers several ways to go about doing that. If your computer doesn't have a CD drive, Network and Hard Disk installs are available. To preconfigure Red Hat Linux to install on multiple, similar computers, you can use the kickstart installation.

Note Whether you are installing Red Hat Linux from the CDs with this book (referred to as Fedora Core) or Red Hat Enterprise Linux CDs directly from Red Hat, Inc., you will use basically the same installer (called anaconda). In this chapter, I step you through the procedure for installing Red Hat Linux from the Fedora CDs. See the section "Installing Red Hat Linux (Fedora or Enterprise)" for information on the few ways that the installation procedure differs for an Enterprise installation.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Quick Installation

It can be a little intimidating to see a thick chapter on installation. But the truth is, if you have a little bit of experience with computers and a computer with common hardware, you can probably install Red Hat Linux pretty easily. The procedure in this section will get you going quickly if you have:

The Red Hat Linux installation CDs (Fedora Core) that come with this book.

A Pentium-class PC (at least 200 MHz for text mode; 400 MHz for GUI) with a built-in, bootable CD-ROM drive, at least 64MB of RAM (for text mode) or 128MB of RAM (for GUI mode). You need at least 510MB of free hard disk space for a Minimum custom install, at least 2.1GB of hard disk space for a personal desktop install, at least 2.6GB of free space for a workstation install, and at least 920MB for a server install. (The Minimum install is configured to be used as a Linux firewall and/or router.) A custom Everything install requires at least 5.8GB of disk space.

For this quick procedure, you must either be dedicating your entire hard disk to Linux, have a preconfigured Linux partition, or have sufficient free space on your hard disk outside any existing Windows partition.

Caution If you are not dedicating your whole hard disk to Red Hat Linux and you don't understand partitioning, skip to the "Detailed Installation Instructions" section in this chapter. That section describes choices for having both Linux and Windows on the same computer.

Here's how you get started:

1. Insert Red Hat Linux (Fedora Core) installation CD #1 into your computer's CD-ROM drive.
2. Reboot your computer.
3. When you see the installation screen, press Enter to begin the installation.

During installation, you are asked questions about your computer hardware and the network connections. After you have completed each answer, click Next. The following list describes the information you will need to enter. (If you need help, all of these topics are explained later in this chapter.)

Media Check — Optionally check each CD to be sure it is not damaged or corrupted.

Language — Choose your language.

Keyboard — Choose your keyboard type.

Mouse — Identify the type of mouse you are using.

Monitor — Identify your monitor model or at least its horizontal and vertical sync.

Upgrade? — If you have an earlier version of Red Hat Linux installed, you can choose Upgrade to upgrade your system without losing data files. Otherwise, you can continue with a new installation.

Install type — Choose a configuration, such as Personal Desktop (for laptop, home, or desktop use), Workstation (desktop plus software development), Server (file, print, Web, and other server software), or Custom (adds selected Linux packages, Minimum, or Everything installs).

Partitions — Either have Red Hat automatically choose your partitions or manually partition yourself (with Disk Druid). With Automatic, you can choose to remove Linux partitions, all partitions, or no partitions (and use existing free space). Because repartitioning can result in lost data, I recommend that you refer to descriptions on repartitioning your hard disk later in this chapter.

Disk Druid — Whether you choose Automatic or Manual partitioning, Disk Druid appears onscreen to let you review or change the partitions.

Boot Loader — Add the GRUB boot manager to control the boot process. (GRUB is described later in this chapter.)

Network configuration — Set up your LAN connection (not dial-up). You can simply choose to get addresses using DHCP, or you can manually enter your computer's IP address, netmask, host name, default gateway, and DNS servers. You can also indicate whether to activate your network when Linux boots.

Firewall — Choose a default firewall configuration. Select High if you plan to connect

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

to the Internet but don't plan to use Linux as a server. Select Medium to allow a select group of services. Choose Custom if you want to offer particular services to the network. Select No Firewall only if you are connected to a trusted network, with no connection to a public network

Language support — Choose to install support for additional languages.

Time zone — Identify the time zone in which you are located.

Root password — Add the root user account password.

Packages — Select to accept the current package list (for the install type you chose) or customize it. For custom installations, choose groups of software packages to install, choose Everything, or Mimimum. (You can also choose separate packages if you like.)

Caution If your computer is connected to the Internet, you should be more selective about which server packages you install because they may pose potential security risks. A misconfigured server can be like an open window to your computer. In a safe environment, however, an Everything install (if you have enough disk space) allows you to follow the procedures in this book without continuously going back and installing new packages from the CD.

Installing packages — To this point, you can quit the install process without having written anything to disk. When you select Next, the disk is formatted (as you chose) and selected packages are installed.

Note After answering the questions, the actual installation of packages takes between 20 and 60 minutes, depending on the number of packages and the speed of the computer hardware. During this time, you will be asked to insert the other Red Hat installation CDs.

Boot disk — Create a boot disk (optional, but recommended).

When installation is done, remove the Red Hat Linux CD and click Exit to reboot your computer. When you see the boot screen, use up and down arrows to select a partition. Linux should boot by default. After Linux boots for the first time, the Red Hat Setup Agent runs to let you set system date and time, sign up for software updates, and install additional CDs. On subsequent reboots, you will see a login prompt. If you need more information than this procedure provides, go to the detailed installation instructions just ahead.

Note If you did a Server or Minimum install where no graphical interface is installed, the Red Hat Setup Agent does not run.

Team LiB

PREVIOUS NEXT

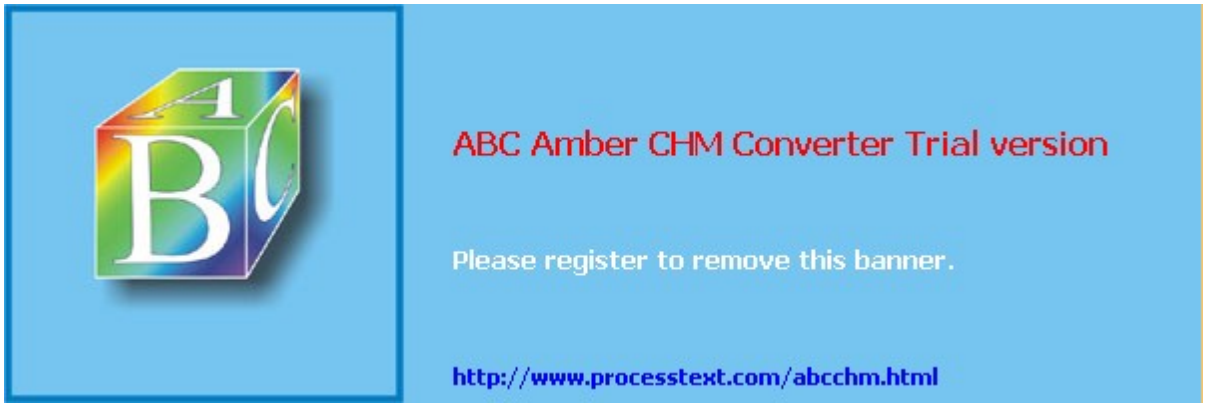

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Detailed Installation Instructions

This section provides more detail on installation. Besides expanding on the installation procedure, this section also provides information on different installation types and on choosing computer hardware.

Tip If anything goes wrong during installation and you get stuck, go to the "Troubleshooting Your Installation" section at the end of this chapter. It gives suggestions for solving common installation problems.

Caution If you are installing a dual-boot system that includes a Windows operating system, try to install the Windows system first and the Red Hat Linux system later. Some Windows systems blow away the Master Boot Record (MBR), making the Red Hat Linux partition inaccessible.

If, when installing Windows or Red Hat Linux, you find that the other operating system is no longer available on your boot screen, don't panic and don't immediately reinstall. You can usually recover from the problem by booting with the Red Hat Linux emergency boot disk, and then using either the grub-install or lilo commands to reinsert the proper MBR. If you are uncomfortable working in emergency mode, seek out an expert to help you.

Installing Red Hat Linux (Fedora or Enterprise)

This chapter details how to install Red Hat Linux from the CDs that come with this book. If you are installing Red Hat Linux from those CDs, you can simply follow the instructions in this chapter. If you are installing Red Hat Enterprise Linux instead, there are a few differences in the installation procedure. Descriptions of the differences between Fedora Core and Red Hat Enterprise Linux are described in Chapter 1. However, here are a few issues you should be aware of if you are using the installation procedure in this chapter to install Red Hat Enterprise Linux.

- Instead of having a 3-CD installation, Red Hat Enterprise Linux consists of a different boot CD for AS and WS installs. After starting installation with the appropriate boot CD, both install types use the same set of additional CDs (marked disc2, disc3, and disc4).
- Installation classes for Fedora and Enterprise are different.
- The names and logos used for Fedora and Enterprise are different.
- Unlike the Fedora installation, which installs all CDs in order, Red Hat Enterprise Linux requires that you insert the boot CD again near the end of the install process.

Besides those differences, an installation of Red Hat Enterprise Linux should match the instructions in this chapter. There are differences in which packages are included with the Fedora and Enterprise distributions, however. (See Appendix B for package descriptions.)

Choosing an installation method

Red Hat Linux offers very flexible ways of installing the operating system. Of course, I recommend installing Red Hat Linux from the CDs that come with this book. However, if you don't have the Red Hat CDs or if you don't have a working CD-ROM drive, you can install Red Hat Linux from any of several different types of media. There are also several special types of installation. The installation types noted here are described fully in the "Special Installation Procedures" section.

Install or upgrade?

First you should determine if you are doing a new install or an upgrade. If you are upgrading an existing Red Hat Linux system to the latest version, the installation process will try to leave your data files and configuration files intact as much as possible. This type of installation takes longer than a new install. A new install will simply erase all data on the Linux partitions (or whole hard disk) that you choose.

From CD, network, or hard disk?

You can install Red Hat Linux from any of the locations described here. Each of the installation procedures from locations other than a CD requires a Red Hat Linux installation boot disk. (Creating an

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

installation boot disk is described later.)

HTTP — Lets you install from a Web page address.

FTP — Lets you install from an FTP site.

NFS image — Allows you to install from any shared directory on another computer on your network using the Network File System (NFS) facility.

Hard disk — If you can place a copy of the Red Hat Linux distribution on your hard disk, you can install it from there. (Presumably, the distribution is on a hard disk partition to which you are *not* installing.)

The following specialty installation type also may be of interest to you:

Kickstart installation — Lets you create a set of answers to the questions Red Hat Linux asks you during installation. This can be a time-saving method if you are installing Red Hat Linux on many computers with similar configurations.

The Red Hat Installation Guide is available from any Red Hat FTP site (such as `ftp.redhat.com`). The location on the `ftp.redhat.com` server for the previous release of Red Hat Linux is:

pub/redhat/linux/9/en/doc/RH-DOCS/rhl-ig-x86-en-9/index.html

Another document you may find useful before installing is the Red Hat Linux Reference Guide (also listed in the `RH-DOCS` directory, as `rhl-rg-en-9.0`). You'll need to check for yourself to find out if Red Hat continues to update the reference guides for Fedora Core.

Choosing computer hardware

This may not really be a choice. You may just have an old PC lying around that you want to try Red Hat Linux on. Or you may have a killer workstation with some extra disk space and want to try out Red Hat Linux on a separate partition or whole disk. To install the PC version of Red Hat Linux successfully (that is, the version on the accompanying CD), the computer must have the following:

x86 processor — Your computer needs an Intel-compatible CPU. With the latest version, Red Hat recommends that you at least have a Pentium-class processor to run Red Hat Linux. For a text-only installation, a 200 MHz Pentium is the minimum, while a 400 MHz Pentium II is the minimum for a GUI installation. Although some 486 machines will work, they cannot be counted on.

Floppy disk drive or CD-ROM — You need to be able to boot up the installation process from either floppy disk or CD-ROM. If you don't have a CD-ROM drive, you need a LAN connection to install Red Hat Linux from a server on the network or figure out a way to copy the contents of the CD to a hard disk.

Hard disk —The minimum amount of space you need varies depending on the installation type and packages you select. If you are an inexperienced user, you want at least 2.1GB of space so you can get the GUI with a personal desktop or workstation install:

- **Personal Desktop** — Requires 2.1GB of disk space.
- **Workstation** — Requires 2.6GB of disk space.
- **Server** — Requires 920MB of disk space.
- **Everything** (Custom) — Requires about 5.8GB.
- **Mimimum** (Custom) — Requires at least 510MB of disk space.

RAM — You should have at least 64MB of RAM to install Red Hat Linux. If you are running in graphical mode, you will probably need at least 128MB. The recommended RAM for GUI mode is 256MB.

Keyboard and monitor — Although this seems obvious, the truth is that you only need a keyboard and monitor during installation. You can operate Red Hat Linux quite well over a LAN using either a shell interface from a network login or an X terminal.

Red Hat Enterprise Linux versions, not included with this book, are available for other architectures, such as Intel Itanium, AMD64, IBM PowerPC, and IBM mainframe. The CDs that come with this book and the installation procedures presented here, however, are specific to PCs. Most of the other software described in this book, however, will work the same in any of those hardware environments. (Check out `http://redhat.com/mirrors` for sites that offer Red Hat Linux for different computer hardware architectures.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The list of hardware supported by Red Hat Linux is available on the Internet at www.redhat.com/hardware.

If your computer's CD-ROM device is connected to a PCMCIA port (such as those that come on laptop computers), you will need to install PCMCIA support during installation. PCMCIA support is available only on Intel-based computers. See the sidebar on installing Red Hat Linux on a laptop for further information.

Beginning the installation

If you feel you have chosen the right type of installation for your needs, you can begin the installation procedure. Throughout most of the procedure, you can click Back to make changes to earlier screens. However, once you are warned that packages are about to be written to hard disk, there's no turning back. Most items that you configure can be changed after Red Hat Linux is installed.

Installing Red Hat Linux on a Laptop

Because laptops can contain non-standard equipment, before you begin installing on a laptop you should find out about other people's experiences installing Linux on your model. Do that by visiting the Linux on Laptops site (`www.linux-on-laptops.com`).

If your laptop doesn't have a built-in CD-ROM drive, you probably need to install from a device connected to a PCMCIA slot on your laptop. PCMCIA slots let you connect a variety of devices to your laptop using credit-card-sized cards (sometimes called PC Cards). Linux supports hundreds of

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

PCMCIA devices. You may use your laptop's PCMCIA slot to install Red Hat Linux from several different types of PCMCIA devices, including:

- A CD-ROM drive
- A LAN adapter

If you would like to know which PCMCIA devices are supported in Linux, see the `SUPPORTED.CARDS` file (located in the `/usr/share/doc/kernel-pcmcia-cs*` directory). In any of these cases, you need the PCMCIA support disk to use the device as an installation medium. The section on creating install disks describes how to create these installation floppy disks. (See Chapter 10 for further information on using Linux on laptops.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution It is quite possible that your entire hard disk is devoted to a Windows 95, 98, 2000, ME, NT, or XP operating system and you may want to keep much of that information after Re

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you are upgrading an existing Red Hat Linux system to this release, you should consider first removing any unwanted packages from your old Red Hat Linux system. Fewer

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

1. **Insert the first CD-ROM in the CD-ROM drive.** If you are installing from a local hard disk or network and you have a floppy disk drive but no CD-ROM drive, you can insert an installation floppy boot disk instead. Refer to the section on creating install disks for information on making the disk (or disks) you need.
2. **Start your computer.** If you see the Red Hat Linux installation screen, continue to the next step.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you don't see the installation screen, your CD-ROM drive may not be bootable. Creating a bootable floppy may be the best way to proceed. However, you may also have the ch

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

3. **Start the boot procedure.** At the boot prompt, press Enter to start the boot procedure in graphical mode. If for some reason your computer will not let you install in graphical mode (16-bit color, 800x600 resolution, framebuffer), refer to the "Choosing Different Install Modes" sidebar. Different modes let you start network installs and nongraphical installs (in case your video card can't be detected).
4. **Media check.** At this point, you may be asked to check your installation media. If so, press Enter to check that the CD is in working order. If one of the CDs is damaged, this step saves you the trouble of getting deep into the install before failing. Repeat this step for each CD in the set; then select Skip to continue.
5. **Continue.** When the welcome screen appears, click Release Notes to see information about this version of Red Hat Linux. Click Next when you're ready to continue.
6. **Choose a language.** When prompted, indicate the language that you would like to use during the installation procedure by moving the arrow keys and selecting Next. (Later, you will be able to add additional languages.) You are asked to choose a keyboard.
7. **Choose a keyboard.** Select the correct keyboard layout (U.S. English, with Generic 101-key PC keyboard by default). Some layouts enable dead keys (on by default). Dead keys let you use characters with special markings (such as circumflexes and umlauts).
8. **Add a mouse.** When prompted, indicate the kind of mouse and click Next. If possible, choose an exact match to the model of mouse you have. Otherwise, you can choose a generic serial mouse (if it connects to a COM port) or a generic PS/2 mouse (if it connects to a PS/2 port). Support also exists for two-button and three-button generic USB mice. For a serial mouse, identify which COM port the mouse is connected to. Generic wheel mice (PS/2 and USB) are also supported. If you are using a two-button mouse, you can click Emulate 3 Buttons. This lets you use the Shift key with a mouse button to emulate the center button from a three-button mouse. (Some USB mice are not well detected during install. Because many USB mice use the same protocol, try selecting Logitech Mouseman Wheel (USB) as the mouse type if you can't match your exact model.)
9. **Select Monitor Configuration.** Scroll down the list to find your monitor's manufacturer; then click the plus sign to choose the model. When you select the model, the correct horizontal and vertical sync rates are added, or you can type your own values. If your model is not found, consult the monitor's manual. Then try a Generic CRT or type the monitor's horizontal and vertical sync rates into the appropriate box. Click Next to continue.
10. **Choose install type.** Select either "Install Red Hat Linux" for a new install or "Upgrade an existing installation" to upgrade an existing version of Red Hat.

11. **Choosing Different Install Modes**

12. Although most computers will enable you to install Red Hat Linux in the default mode (graphical), there may be times when your video card does not support that mode. Also, though the install process will detect most computer hardware, there may be times when your hard disk, Ethernet card, or other critical piece of hardware cannot be detected and will require you to enter special information at boot time.
13. The following is a list of different installation modes you can use to start the Red Hat Linux install process. You would typically only try these modes if the default mode failed (that is, if the screen was garbled or hardware wasn't detected). For a list of other supported modes, refer to the `/usr/share/doc/anaconda*/command-line.txt` file or press F2 to see short descriptions of some of these types.

    linux text: Type **linux text** to run installation in a text-based mode. Do this if installation doesn't seem to recognize your graphics card. The installation screens aren't as pretty, but they work just as well.
    linux lowres: Type **linux lowres** to run installation in 640x480 screen resolution for graphics cards that can't support the higher resolution.
    linux nofb: Type **linux nofb** to turn off frame buffer.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

linux noprobe: Normally, the installation process will try to determine what hardware you have on your computer. In `noprobe` mode, installation will not probe to determine your hardware; you will be asked to load any special drivers that might be needed to install it.

linux mediacheck: Type **linux mediacheck** to check your CDs before installing. Because media checking is done next in the normal installation process, you should do this only to test the media on a computer you are not installing on.

linux rescue: The **linux rescue** mode is not really an installation mode. This mode boots from CD, mounts your hard disk, and lets you access useful utilities to correct problems preventing your Linux system from operating properly.

linux vnc vncpassword=****:** Run your computer as a vnc server to allow other computers to install Red Hat Linux from your computer. The password (represented by asterisks) must be at least six characters. The client must connect to display :1 and supply the password to be able to install.

linux dd: Type **linux dd** if you have a driver disk you want to use to install.

linux expert: Type **linux expert** if you believe that the installation process is not properly auto-probing your hardware. This mode bypasses probing so you can choose your mouse, video memory, and other values that would otherwise be chosen for you.

linux askmethod: Type **linux askmethod** to have the installation process ask where to install from (local CD, NFS image, FTP, HTTP, or hard disk).

linux updates: Type **linux updates** to install from an update disk.

You can add other options to the linux boot command to identify particular hardware that is not being detected properly. For example, to specify the number of cylinders, heads, and sectors for your hard disk (if you believe the boot process is not detecting these values properly), you could pass the information to the kernel as follows: `linux hd=720,32,64`. In this example, the kernel is told that the hard disk has 720 cylinders, 32 heads, and 64 sectors. You can find this information in documentation that comes with your hard disk (or stamped on the hard disk itself on a sticker near the serial number).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note To upgrade, you must have at least a Linux 2.0 kernel installed. With an upgrade, all of your configuration files are saved as *filename*.rpmsave (for example, the hosts file is saved as hosts.rp

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Warning

The personal desktop and workstation installation types do not install server packages or many system administration tools. To use most of the administration and server features described

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

For a new install, you must choose one of the following types (also referred to as classes) of installation. For any of these installation types, you will have the opportunity to install a set of preset packages or customize that set.

Personal Desktop — Installs software appropriate for a home or office personal computer or laptop computer. This includes the GNOME desktop (no KDE) and various desktop-related tools (word processors, Internet tools, and so on). Server tools, software development tools, and many system administration tools are not installed.

Workstation — Similar to a Personal Desktop installation but adds tools for system administration and software development. (Server software is not installed.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

Any Linux partitions or free space on your hard disk(s) will be assigned to the new installation with the Personal Desktop or Workstation types of installation. Any Windows partition

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Server — Server installs the software packages that you would typically need for a Linux server (in particular, Web server, file server, and print server). It does not include many other server types (DHCP, mail, DNS, FTP, SQL, or news servers). The default server install does not include a GUI (so you'd better know how to use the shell). This install type also erases all hard disks and assigns them to Linux by default.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

This is a big one. In case you didn't catch the previous paragraph, Server installs erase the whole hard disk by default! If you have an existing Windows partition that you want to ke

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Custom System — You are given the choice of configuring your own partitions and selecting your own software packages.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note

With an Everything custom install, you will have all the desktop, server, and development tools that come with Red Hat Linux. If you have the disk space, an Everything install saves yo

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

At this point, the procedure will continue through a Custom System installation. Even though different install classes choose different partitioning methods by default, in all cases you have the choice to see and change the partitioning that was chosen for you.

14. **Choose your partitioning strategy.** You have two choices related to how your disk is partitioned for Red Hat Linux installation:

    Automatically partition— With this selection, all Linux partitions on all hard disks are erased and used for the installation. The installation process automatically handles the partitioning. (It does give you a chance to review your partitioning, however.)

    Manually partition with Disk Druid — With this selection, the Disk Druid utility is run to let you partition your hard disk.

    Click Next to continue.

15. **Choose partitioning.** If you selected to have the installer automatically partition for you, you can choose from the following options:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you selected to use Disk Druid for partitioning, refer to the section on partitioning your hard disk later in this chapter for details on using those tools.

Remove all Linux partitions on this system — Windows and other nonLinux partitions remain intact with this selection.

Remove all partitions on this system — This erases the entire hard disk.

Keep all partitions and use existing free space — This only works if you have enough free space on your hard disk that is not currently

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

assigned to any partition.

If you have multiple hard disks, you can select which of those disks should be used for your Red Hat Linux installation. Turn the Review check box on to see how Linux is choosing to partition your hard disk. Click Next to continue.

After reviewing the Partitions screen, you can change any of the partitions you choose, providing you have at least one root (/) partition that can hold the entire installation and one swap partition. A small `/boot` partition (about 100MB) is also recommended. The swap partition is usually twice the size of the amount of RAM on your computer (for example, for 128MB RAM you could use 256MB of swap). Having less than double your physical memory for swap can cause performance problems.

Click the Next button (and select OK to accept any changes) to continue.

16. **Configure boot loader.** All bootable partitions and default boot loader options are displayed. By default, the install process will use the GRUB boot loader, install the boot loader in the master boot record of the computer, and choose Red Hat Linux as your default operating system to boot.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you keep the GRUB bootloader, you have the option of adding a GRUB password. The password protects your system from having potentially dangerous kernel options sent to the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

17. The names shown for each bootable partition will appear on the boot loader screen when the system starts. Change a partition name by clicking it and selecting Edit. To change the location of the boot loader, click "Configure advanced boot loader options" and continue to the next step. (If the defaults are okay, skip the next step.)
18. To choose where to store the boot loader, select one of the following:

    Master Boot Record (MBR) — This is the preferred place for GRUB. It causes either GRUB or LILO to control the boot process for all operating systems installed on the hard disk.

    First Sector of Boot Partition — If another boot loader is being used on your computer, you can have GRUB installed on your Linux partition (first sector). This lets you have the other boot loader refer to your GRUB boot loader to boot Red Hat Linux.

    You can choose to add Kernel Parameters (which may be needed if your computer can't detect certain hardware). You can select to use linear mode (which was once required to boot from a partition on the disk that is above cylinder 1024, but is now rarely needed).

Note For more information on GRUB, refer to the section on bootloaders later in this chapter.

19. **Configure networking.** At this point, you are asked to configure your networking. This applies only to configuring a local area network**.** If you will use only dial-up networking, skip this section by clicking Next. If your computer is not yet connected to a LAN, you should skip this section.

    Network address information is assigned to your computer in two basic ways: statically (you type it) or dynamically (a DHCP server provides that information from the network at boot time). One Network Device appears for each network card you have installed on your computer. The first interface is `eth0`, the second is `eth1`, and so on. Repeat the setup for each card by selecting each card and clicking Edit.

Cross-Reference Refer to Chapter 15 for descriptions of IP addresses, netmasks, and other information you need to set up your LAN and to Chapter 16 for information related to domain names.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

With the "Edit Interface eth0" window displayed, add the following:

Configure using DHCP — If your IP address is assigned automatically from a DHCP server, a checkmark should appear here. With DHCP checked, you don't have to set other values on this page. Remove the checkmark to set your own IP address.

IP Address — If you set your own IP address, this is the four-part, dot-separated number that represents your computer to the network. How IP addresses are formed and how you choose them is more than can be said in a few sentences (see Chapter 15 for a more complete description). An example of a private IP address is 192.168.0.1.

Netmask — The netmask is used to determine what part of an IP address represents the network and what part represents a particular host computer. An example of a netmask for a Class C network is 255.255.255.0.

Activate on boot — Besides this information, you should indicate whether you want the network to start at boot time (you probably do if you have a LAN).

Click OK. Then add the following information on the main screen:

Set the hostname — This is the name identifying your computer within your domain. For example, if your computer were named "baskets" in the `handsonhistory.com` domain, your full hostname may be `baskets.handsonhistory.com`. You can either set the domain name yourself (manually) or have it assigned automatically, if that information is being assigned by a DHCP server (automatically via DHCP).

Gateway — This is the IP number of the computer that acts as a gateway to networks outside your LAN. This typically represents a host computer or router that routes packets between your LAN and the Internet.

Primary DNS — This is the IP address of the host that translates computer names you request into IP addresses. It is referred to as a Domain Name System (DNS) server. You may also have Secondary and Tertiary name servers in case the first one can't be reached. (Most ISPs will give you two DNS server addresses.)

Cross-Reference To configure your LAN after installation, see Chapter 15.

20. **Choose a firewall configuration.** The use of a firewall has significant impact on the security of your computer. If you are connected to the Internet or to another public network, a firewall can limit the ways an intruder may break into your Linux system. Here are your choices for configuring a firewall during installation:

High — Select this security level if you are connecting your Linux system to the Internet for Web browsing and file downloading (FTP) but don't plan to offer your system as a server to the Internet. Only explicitly defined connections are accepted. To allow Web browsing and basic network setup, DNS replies and DHCP (to serve addresses) are allowed at this level.

Medium — Select this security level if you want to block access to ports used to offer most basic TCP/IP services (standard, reserved ports lower than 1023). This selection denies access to ports used for NFS servers, remote X clients, and the X font server. Those services are available primarily to other computers your own local LAN.

No firewall — Select this security level if you are not connected to a public network and do not want to deny requests for services from any computer on your local network. Of course, you can still restrict access to services by starting up only the services you want to offer and by using configuration files to restrict access to individual services.

If you know you want to allow access to particular services, you can click Customize and allow incoming requests for the following: DHCP, SSH, Telnet, WWW, Mail, and/or FTP services. You can also add a comma-separated list of port numbers to the Other Ports box

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

to open access to those ports. (The `/etc/services` file lists which services are associated with which port numbers.)
If you have a LAN that consists of trusted computers, you can click the box representing your interface to that LAN (probably `eth0`). Clicking the box allows access to any services you care to share with the computers on your LAN.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Adding firewall rules here results in rules being added to the /etc/sysconfig/iptables file. The rules are run from the /etc/init.d/iptables start-up script when you boot your system. To

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

21. **Choose language support.** Your installation language should be selected automatically as your default language on this screen. You can select to install support for additional languages by clicking the check boxes next to the languages you want. You can click the Select All button to install all supported languages to your system.
22. **Choose a time zone.** Select the time zone from the list of time zones shown. Either click a spot on the map or choose from the scrolling list. To see a more specific view of your location, click World and choose your continent. From the UTC Offset tab, you can choose a time zone according to the number of hours away from Greenwich Mean Time (GMT), known as the UTC offset.
23. **Set root password.** You must choose a password for your root user at this point. The root password provides complete control of your Red Hat Linux system. Without it, and before you add other users, you will have no access to your own system. Enter the Root Password, and then type it again in the Confirm box to confirm it. (Remember the root user's password and keep it confidential! Don't lose it!) Click Next to continue.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Use the passwd command to change your password later. See Chapter 14 for suggestions on how to choose a good password. See Chapter 11 for information on setting up user

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

24. **Select Packages.** You are presented with groups of packages at this point. Which packages are selected by default depends on the type of installation you chose earlier. In general, either more workstation-oriented or server-oriented packages are selected. Select the ones you want and click Next.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip You can override your package selections by choosing "Minimal" or "Everything" install groups. Disk space requirements for those install types are described earlier in this chapt

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

25. Because each group represents several packages, you can click the Details button next to each group to select more specifically the packages within that group. Because workstation and personal desktop installations don't add any server packages, this is a good opportunity to add server packages for the services you expect to use.

Cross-Reference Appendix B describes the software packages that come with Red Hat Linux.

26.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip A listing of all of the software packages is contained in the file RedHat/base/comps.xml on the first Red Hat installation CD-ROM. It's in XML format, so it is incomprehensible if you

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

27. **About to Install.** A screen tells you that you are about to begin writing to hard disk. You can still back out now, and the disk will not have changed. Click Next to proceed. (To quit without changes, eject the CD and restart the computer.) Now the file systems are created and the packages are installed. This typically takes from 20 to 60 minutes to complete, although it can take much longer on older computers.

    You are prompted to insert additional installation CDs as they are needed.
28. **Create boot disk.** If you want to create a boot disk, insert a blank floppy and click Next to create the boot floppy. I recommend that you do this if you have a floppy disk drive, and label the disk with the name of the computer, operating system, and release number. (You can still skip it by clicking No.)
29. **Finish installing.** When you see the "Congratulations" screen, you are done. Note the links to Red Hat Linux information, eject the CD and click Exit.

Your computer will restart. If you installed GRUB, you will see a graphical boot screen that displays the bootable partitions. Press the Up Arrow or Down Arrow keys to choose the partition you want to boot, and press Enter. If Linux is the default partition, you can simply wait a few moments and it will boot automatically.
The first time your system boots after installation, the Red Hat Setup Agent runs to do some initial configuration of your system. The next section tells how Red Hat Setup Agent works.

Running Red Hat Setup Agent

The first time you boot Red Hat Linux, after it is installed, the Red Hat Setup Agent runs to configure some initial settings for your computer.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The Red Hat Setup Agent only runs automatically if you have configured Red Hat Linux to boot to a graphical login prompt. To start it from a text login, log in as root and switch to init state 5 temporarily (type **init 5**). Log in to the graphical prompt. From a Terminal window as root user, type: # **rm /etc/a**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The first screen you see is the Welcome screen. From there, click the Next button to step through each procedure as follows:

Date and Time Configuration — You can manually enter the date (click on the calendar) and time (select hour, minutes, and seconds) or use the network time protocol to have your date and time set automatically from a known time server. Click on Enable Network Time Protocol (NTP), then select a time server by clicking on the down arrow and selecting a site. Network Time Protocol (NTP) is a service that allows computers to synchronize their date and time clocks with reliable time servers.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Red Hat offers two time servers you can use (click the down arrow in the server box to see them). Or you can type in your own time server. It is better to type an IP address than a

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Setting NTP in this way adds your chosen NTP server to the `/etc/ntp.conf` file (see the `server` option). To check that time has been synchronized, type the `ntptrace` command. You should not have to change your firewall for NTP to work, because NTP attempts to punch a hole through your firewall to synchronize your time.

User Account — For your daily use of Red Hat Linux, you should have your own user account. You should typically log in with this user name (of your choosing) and use only the root user to perform administrative tasks. In the first of the four text boxes on the screen, type a user name (something like `jparker` or `alanb`). Next, type your full name (like John W. Parker or Alan Bourne). Then type your password in the Password box and again in the Confirm Password box. Click Forward.

If some form of network authentication is used, such as LDAP, Kerberos, or SMB authentication, you can click the Use Network Login button. See the "Enabling Authentication" sidebar for information on choosing different authentication types.

Enabling Authentication

In most situations, you will enable shadow passwords and MD5 passwords (as selected by default) to authenticate users who log in to your computer from local passwd and shadow password files. To change that behavior, you can select the Use Network Login button during the User Account setup during Red Hat Setup Agent (firstboot).

The shadow password file prevents access to encrypted passwords. MD5 is an algorithm used to encrypt passwords in Linux and other UNIX systems. It replaces an algorithm called crypt, which was used with early UNIX systems. When you enable MD5 passwords, your users can have longer passwords that are harder to break than those encrypted with crypt.

If you are on a network that supports one of several different forms of network-wide authentication, you may choose one of the following features:

Configure NIS. Select this button and type the NIS Domain name and NIS server location if your network is configured to use the Network Information System (NIS). Instead of selecting an NIS Server, you can click the checkbox to broadcast to find the server on your network. For more information on NIS, see Chapter 23.

Configure LDAP. If your organization gathers information about users, you can click this button to search for authentication information in an LDAP server. You can enter the LDAP Server name and optionally an LDAP distinguished name to look up the user information your system needs.

Configure Hesiod. If your organization uses Hesiod for holding user and group informnation in DNS, you can add the LHS (domain prefix) and RHS (Hesiod default domain) to use for doing Hesiod queries.

Configure Kerberos Support. Click this button to enable network authentication services available through Kerberos. After enabling Kerberos, you can add information about a Kerberos Realm (a group of Kerberos servers and clients), KDC (a computer that issues Kerberos tickets), and Admin server (a server running the Kerberos kadmind daemon).

Configure SMB. Click this tab to configure your computer to use Samba for file and print sharing with Windows systems. If you enable SMB authentication, you can enter the name of the SMB server for your LAN and indicate the Workgroup you want your computer to belong to.

Sound Card — The Setup Agent tries to detect your sound card. Click on Play test sound, to check if the card is working. If you hear the sound, click Yes. Then click Next to continue.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Install Additional Software — If you have Red Hat Linux CDs other than those that come with this book, you can install them now by inserting the CD you want to install and clicking the appropriate button.

The Red Hat Setup Agent is complete. See Chapter 3 for a description of how to log in to Red Hat Linux and start learning how to use Linux.

When Red Hat Linux starts up the next time, it will boot up normally to a login prompt. For this release of Red Hat Linux, a new graphical boot screen is displayed (instead of a scrolling list of services starting up). If you miss the old scrolling list, you can view it by clicking the Details button or by pressing Ctrl+Alt+F1.Then go back again by pressing Ctrl+Alt+F8.

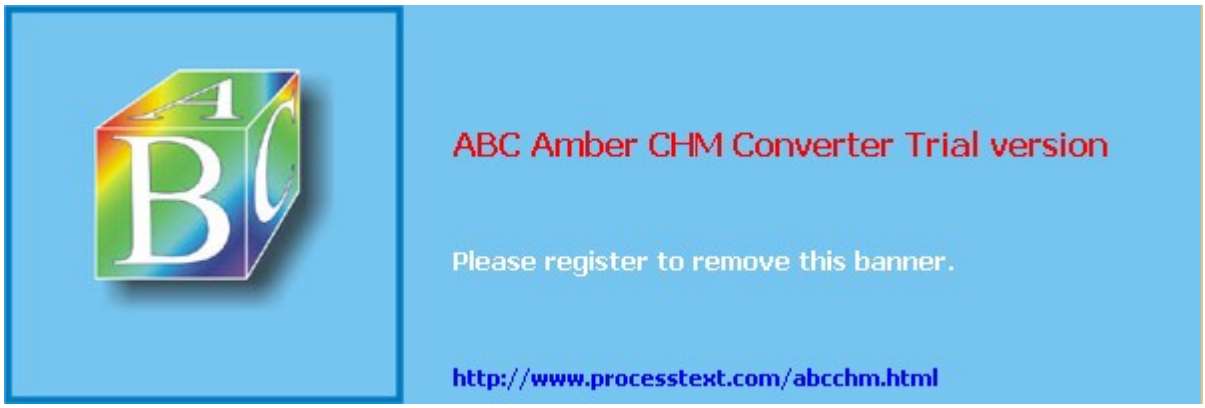

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Installing More Red Hat Linux Packages

Unless you did an "Everything" install, as you go through this book you will probably find that you want to add some Red Hat Linux software packages later. To do that, you can use `rpm` (described in Chapter 5), a general-purpose command for installing any software packages in RPM format. Or you can use the Red Hat Add and Remove Software window (redhat-config-packages), which provides the same interface for installing packages that was used in the original installation procedure.

To open the Red Hat Add and Remove Software window, select System Settings → Add/Remove Applications from the main menu on the desktop (or type **redhat-config-packages** from a Terminal window). The window appears as shown in Figure 2-1.

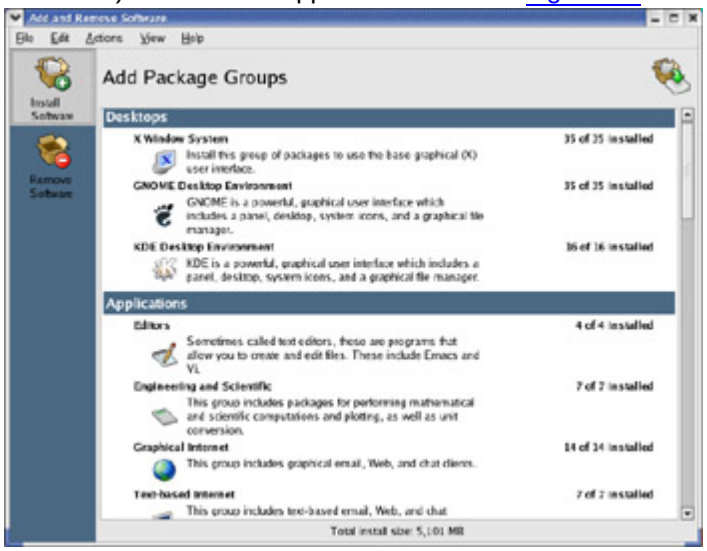

Figure 2-1: Change software packages after Red Hat installation using the Add and Remove Software window.

Scroll through the major package groups using the scroll bar. After the name and description of each group there is a set of numbers in brackets that shows how many packages are in the group and how many of those are actually installed. Click Details to see the listing of all packages in the group.

The details of each group shows the standard packages in each group (ones that are required if the group is installed at all) and ones that are extra. To add or remove packages from your system, click the check box next to the package so that it is on or off, respectively. Repeat this for each package group, and close the details window.

When you have made all your add or delete selections, click Update on the main Add and Remove Software window. Insert each CD as it is requested. Packages will be added or removed, as you requested.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Special Installation Procedures

If you don't want to, or can't, use the procedure to install Red Hat Linux from CD, the procedures in this section give you alternatives. The first procedure tells how to install Red Hat Linux from alternative media (using FTP, HTTP, NFS, or hard disk installs). The next procedure describes how to do kickstart installations.

Installing from other media

Your CD-ROM drive may not be working, you may not have a CD-ROM drive, or maybe you just want to install a new version of Red Hat Linux that is available from another computer on the network. In any of these cases, Red Hat will let you install Linux from a Web server (HTTP), an FTP server, a shared NFS directory, or local hard disk.

Note To use HTTP, FTP, or NFS installations, your computer must be connected to a LAN that can reach the computer containing the Red Hat Linux distribution. You cannot use a direct dial-up connection. For a local hard disk install, the distribution must have been copied to a local disk that is not being used for installation. See the section "Setting up an HTTP, FTP, or NFS install server" for details on copying the distribution and making it available.

Beginning installation

To begin any of these types of installation, you need to boot the installation process from a floppy disk. Refer to the section on creating boot disks later in this chapter. You also need to know the location of an install server — a Red Hat Linux system on your network that has the Red Hat Linux distribution available.

With the boot disk in hand (`bootdisk.img` and possibly `drvnet.img` for network install, or `pcmciadd.img` for a laptop install), start installation for these install types as follows:

Note With this version of Red Hat Linux, you can also use CD #1 that comes with this book to start a network or hard disk install. Instead of using a floppy disk, as described below, you can insert Red Hat Linux installation CD #1, reboot the computer, then type **linux askmethod**. You can then proceed with the steps shown below.

1. **Insert the floppy boot disk.**
2. **Reboot the computer.** You should see the "Welcome to Red Hat" boot screen.
3. **Press Enter.** You are prompted to select a language.
4. **Select the language.** You are prompted to choose a keyboard type.
5. **Select your keyboard type.** You are prompted to select an installation method.
6. **Choose the installation method.** Select any of the following installation methods: NFS image, FTP, HTTP, or Hard drive.
7. **Configure the network card.** For any of the network installs, you are asked to select your Ethernet card from the list shown. (This may be detected automatically.) If your card is not on the list, you need to create a floppy from the `drvnet.img` file, containing drivers for other network cards. (To use the `drvnet.img` file, press F2, insert the floppy containing the image, and click OK. Choose your card from the list and click OK again.)
8. **Configure TCP/IP.** For any of the network install types (NFS, FTP, and HTTP), you are prompted to configure TCP/IP for your computer. (See the section on configuring networking earlier in this chapter for information on how to add to these fields.)
9. **Identify the location of the Red Hat Linux distribution.** You must identify the NFS server name, FTP site name, or Web site name that contains the Red Hat directory that holds the distribution. Or, if you are installing from hard disk, you must identify the partition containing the distribution and the directory that actually contains the Red Hat directory.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note For an FTP install, if you are not downloading from an anonymous FTP site, you must select "Use non-anonymous FTP" when you identify the server and directory. You will need a user name and password that has access to the shared directory.

10. **Do a checksum.** You are asked if you want to do a checksum on the first disk image (.iso file of the first installation CD) found in that directory. Select Test if you do. This will verify that the image is not corrupted before you begin. Repeat this for each disk image.
11. **Continue with installation.** If the distribution is found in the location you indicated, continue the installation as described in the previous section. The install is text-based.

The next section describes how to set up your own server for installing Red Hat Linux.

Setting up an HTTP, FTP, or NFS install server

If you have a LAN connection from your computer to a computer that has about 2.5GB of disk space and offers NFS, FTP, or Web services, you can install Red Hat Linux from that server. Likewise, you can install from a spare disk partition by using a hard disk install. The following procedures let you set up a Linux install server by either copying all files from the three CDs or by copying images of the three CDs.

Configuring an install server using files

To do an FTP or HTTP install, you must copy the files from the installation CDs to a directory that you make available to the network. Because there are three installation CDs in the Red Hat Linux distribution, you can't just identify the location of a mounted CD. You must install the contents of all three CDs in the same directory structure on the server's hard disk. For example, you could do the following:

mkdir /tmp/rh

mount /mnt/cdrom *With first CD inserted*

cp -r /mnt/cdrom/* /tmp/rh/

umount /mnt/cdrom ; eject /mnt/cdrom

mount /mnt/cdrom *With second CD inserted*

cp -r /mnt/cdrom/* /tmp/rh/

umount /mnt/cdrom ; eject /mnt/cdrom

mount /mnt/cdrom *With third CD inserted*

cp -r /mnt/cdrom/* /tmp/rh/

Just type **y** when it asks to overwrite some files. The distribution directory must contain at least the RPMS and base directories, which must include all necessary software packages. In this example, all files were copied. Setting up an NFS install server requires copying CD images to the shared NFS directory.

Configuring an install server using disk images

Instead of copying all files from the three installation CDs, you can copy the entire images of each of the three disks to your hard disk for NFS or hard disk installs. To do this, insert the first installation CD and type the following:

mkdir /tmp/rh

dd if=/dev/cdrom of=/tmp/rh/disk1.iso *With first CD inserted*

umount /mnt/cdrom ; eject /mnt/cdrom

dd if=/dev/cdrom of=/tmp/rh/disk2.iso *With second CD inserted*

umount /mnt/cdrom ; eject /mnt/cdrom

dd if=/dev/cdrom of=/tmp/rh/disk3.iso *With third CD inserted*

umount /mnt/cdrom ; eject /mnt/cdrom

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

NFS server

Add an entry to the `/etc/exports` file to share the distribution directory you created. Remember, that for NFS installs, this directory must contain CD ISO images. The following entry would make the directory available in read-only form to any computer:

/tmp/rh (ro)

Next, restart NFS by typing the following as root user:

/etc/init.d/nfs restart

To set the NFS service to be on permanently (it is off by default), type the following as root:

chkconfig nfs on

Web server

If your computer is configured as a Web server, you need to simply make the distribution directory available. For example, after creating the distribution directory as described above, type the following:

ln -s /tmp/rh /var/www/html/rh

If your computer were named `pine.handsonhistory.com`, you would identify the install server as `pine.handsonhistory.com` and the directory as `rh`.

FTP server

If your computer is configured as an FTP server, you need to make the distribution directory available in much the same way you did with the Web server. For example, after creating the distribution directory as described above, type the following:

ln -s /tmp/rh /var/ftp/pub/rh

If your computer were named `pine.handsonhistory.com`, you would identify the install server as `pine.handsonhistory.com` and the directory as `pub/rh`.

Hard disk install

With the ISO images of each CD copied to a disk partition that is not being used for your Red Hat Linux install, you can use the hard disk install. If the ISO images exist in the `/tmp/rh` directory of the first partition of your IDE hard disk, you could identify the device as `/dev/hda1` and the directory holding the images as `/tmp/rh`.

Performing a kickstart installation

If you are installing Red Hat Linux on multiple computers, you can save yourself some trouble by preconfiguring the answers to questions asked during installation. The method of automating the installation process is referred to as a `kickstart` installation.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

Based on the information you provide in your urks.cfg file, kickstart will silently go through and install Red Hat Linux without intervention. If this file is not correct, you could easily remove ve

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The general steps of performing a kickstart installation are as follows:

1. **Create a kickstart file.** The kickstart file, named `ks.cfg`, contains the responses to questions that are fed to the installation process.
2. **Install kickstart file.** You have to place the `ks.cfg` on a floppy disk, on a local hard disk, or in an accessible location on the network.
3. **Start kickstart installation.** When you boot the installation procedure, you need to identify the location of the `ks.cfg` file.

Creating the kickstart file

A good way to begin creating your kickstart file is from a sample `ks.cfg` file. When you install Red Hat Linux, the installation process places a file called `anaconda-ks.cfg` into the `/root` directory. You can use this file as the basis for the `ks.cfg` file that you will use for your kickstart installs.

The particular `/root/anaconda-ks.cfg` file you get is based on the information you entered during a regular installation (CD, NFS, and so on). Presumably, if you are installing Red Hat Linux on other computers for the same organization, multiple computers may have a lot of the same hardware and configuration information. That makes this a great file for you to start creating your `ks.cfg` file from.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note For further details about how to use kickstart, refer to the Red Hat Linux Configuration Guide. You can get this guide from any Red Hat mirror site. To use a more graphical tool for config

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To start, log in as the root user. Then make a copy of the `anaconda-ks.cfg` file to work on.

cp anaconda-ks.cfg ks.cfg

Use any text editor to edit the `ks.cfg` file. Remember that required items should be in order and that any time you omit an item, the user will be prompted for an answer. Entries from a `ks.cfg` file that was created from a regular CD installation of Red Hat Linux are used as a model for the descriptions below. You should start with your own `anaconda-ks.cfg` file, and as a result, your file will start out somewhat differently. Commented lines begin with a pound sign (#).

The first line in the `ks.cfg` file should indicate whether the installation is an upgrade or an install. The `install` option runs a new installation. You can use the `upgrade` keyword instead to upgrade an existing system. (For an upgrade, the only requirements are a language, an install method, an install device, a keyboard, and a boot loader.)

install

The required `lang` command sets the language (and to be more specific, the country as well) in which Red Hat is installed. The value is U.S. English (`en_US.UTF-8`) by default.

lang en_US.UTF-8

You can install multiple languages to be supported in Red Hat Linux. Here is an example of the default being set to United States English:

langsupport –default en_US..UTF-8 en_US.UTF-8

The required `keyboard` command identifies a United States (`us`) keyboard by default. More than 70 other keyboard types are supported.

keyboard us

The required `mouse` command identifies the mouse type. The following example shows a generic three-button PS/2 type mouse (`generic3ps/2`). You could replace it with a two-button serial mouse (`generic`), a two-button PS/2 mouse (`genericps/2`), or a Microsoft IntelliMouse (`msintellips/2`). For a serial mouse, identify the device (for example, `--device ttyS0` for the COM1 port). To configure other types of mice, you must run `mouseconfig`. To see other mouse types, run the `mouseconfig --help` command.

mouse generic3ps/2 –device psaux

The optional `xconfig` command can be used to configure your monitor and video card. If you use the `skipx` command instead (as shown in the following code sample), no X configuration is done. (After the system is installed, run `redhat-config-desktop` to set up your X configuration.) When you use the `xconfig` command, you can identify the type of X server to use based on your video card (`--card`) and monitor specs (`--hysnyc` and `--vsync`). A handful of other options enable you to set the color depth in bits (`--depth`), the screen resolution (`--resolution`), whether the default desktop is GNOME or KDE (`--defaultdesktop`), whether the login screen is graphical (`--startxonboot`), and the amount of RAM on your video card (`--videoram`). (All the information after `xconfig` should actually appear on one line.)

skipx

or

xconfig --card "ATI Mach64 3D Rage IIC" --videoram 2048 --hsync 30-95

--vsync 50-180 --resolution 800x600 --depth 16 --startxonboot

defaultdesktop gnome

The optional `network` command lets you configure your Red Hat system's interface to your network. The example tells your computer to get its IP address and related network information from a DHCP server (`--bootproto dhcp`). If you want to assign a particular IP address, use the `--bootproto static` option. Then change the IP address (`--ip`), netmask (`--netmask`), IP address of the gateway (`--gateway`), and IP address of the DNS server (`--nameserver`) to suit your system. You can also add a hostname (`--hostname`).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Although the network values appear to be on three lines, all values must be on the same line.

network --bootproto dhcp

or

network --device eth0 --bootproto static --ip 192.168.0.1 --netmask 255.255.254.0 --gateway 192.168.0.1 --nameserver 192.168.0.254-–hostname duck.ab.com

The `rootpw` command sets the password to whatever word follows (in the example, `paSSword`). It is a security risk to leave this password hanging around, so you should change this password (with the `passwd` command) after Linux is installed. You also have the option of adding an encrypted password instead (`--iscrypted g.UJ.RQeOV3Bg –enablemd5`).

rootpw paSSword

or

rootpw --iscrypted g.UJ.RQeOV3Bg –-enablemd5

The `firewall` command lets you set the default firewall used by your Red Hat Linux system. The default value is `medium` (medium security). You can also set `firewall` to `high` (high security) or `disabled` (no firewall). (These values are described in the installation procedure earlier in this chapter.) As you can see in the example, you can optionally indicate that there be no restrictions from host computers on a particular interface (`--trust eth0`). You can also allow an individual service (`--ssh`) or a particular *`port`*`:`*`protocol`* pair (`--port 1234:upd`).

firewall --high --trust eth0 --ssh --port 1234:udp firewall --high

The required `authconfig` command sets the type of authentication used to protect your user passwords. The `--enableshadow` option enables the `/etc/shadow` file to store your passwords. The `--enablemd5` option enables up to 256 character passwords. (You would typically use both.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

authconfig --enableshadow --enablemd5

The `timezone` command sets the time zone for your Linux system. The default, shown here, is United States, New York (`America/New_York`). The `–utc` option indicates that the computer's hardware clock is set to UTC time. If you don't set a time zone, `US/Eastern` is used. Run the `timeconfig` command to see other valid time zones.

timezone --utc America/New_York

The `bootloader` command sets the location of the boot loader (GRUB, by default). For example, `--location=mbr` adds GRUB to the master boot record. (Use `--location= none` to not add GRUB.) You can also add kernel options to be read at boot time using the `append` option (`--append hdd=ide-scsi`) or an optional password for GRUB (`--password=`*`GRUBpassword`*).

bootloader --location=mbr

password=GRUBpassword

Partitioning is required for a new install, optional for an upgrade. The code that follows is from the sample `ks.cfg` file. The `clearpart --linux` value removes existing Linux partitions (or use `--all` to clear all partitions) on the first hard drive (`--drives=hda`). The `part /boot`, `/` and `swap`, sets the file system type (`--fstype`) and partition name (`onpart`) for each partition assignment. You can also set sizes of the partitions (`--size`) to however many MB you want.

The following is the partition information you requested

Note that any partitions you deleted are not expressed

here so unless you clear all partitions first, this is

not guaranteed to work

#clearpart --linux --drives=hda

#part /boot --fstype ext3 --size=100 --ondisk=hda

#part / --fstype ext3 --size=700 --grow --ondisk=hda

#part swap --size=128 --grow --maxsize=256 --ondisk=hda

To indicate which packages to install, begin a section with the `%packages` command. (A few examples follow.) Designate whole installation groups, individual groups, or individual packages. After `%packages`, start an entry with an `@` sign for a group of packages, and add each individual package by placing its name on a line by itself. Here is an example:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note You can find a listing of package groups and individual packages on the Red Hat installation CD. Find the comps file in the RedHat/base directory. However, if you start with the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
%packages
@ Administration Tools
@ Authoring and Publishing
@ Core
@ DNS Name Server
@ Development Libraries
@ Development Tools
    .
    .
    .
@ Web Server
@ Windows File Server
@ X Window System
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The %packages command is not supported for upgrades. To do an Everything install, you can remove the package names shown. Then, after the %packages line, you can

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The %post command starts the post-installation section. After it, you can add any shell commands you want to run after installation is completed. By default, you should have useradd commands for users you added during installation. You can also use the usermod command to add the user's password.

%post

/usr/sbin/useradd jake

chfn –f 'John W. Jones' jake

/usr/sbin/usermod –p '1ðrãæàÕÖà$kQUMYbFhOh79wECxnTuaH.' jake
At this point you should have a working ks.cfg file.

Installing the kickstart file

Once the ks.cfg file is created, you need to put it somewhere accessible to the computer doing the installation. Typically, you will place the file on a floppy disk. However, you can also put the file on a computer that is reachable on the network or on a hard disk.

To copy the file to a floppy disk, create a DOS floppy and copy the file as follows:

mcopy ks.cfg a:

When you do the Red Hat Linux kickstart installation, have this floppy disk with you.
Being able to place the ks.cfg file on a computer on the network requires a bit more configuration. The network must have a DHCP or a BOOTP server configured that is set up to provide network information to the new install computer. The NFS server containing the ks.cfg file must export the file so that it is accessible to the computer installing Linux. To use a ks.cfg file from the local hard disk, you can place the file on any partition that is a Windows (VFAT) or Linux (ext3) partition.

Booting a kickstart installation

If the kickstart file (ks.cfg) has been created and installed in an accessible location, you can start the kickstart installation. Here is an example of how you can do a kickstart installation using the Red Hat Linux CD and a floppy containing a ks.cfg file:

1. Insert the Red Hat Linux CD and restart the computer.
2. When you see the boot prompt, insert the floppy containing the ks.cfg file and type the following (quickly, before the installation boots on its own):
3. boot: **linux ks=file:fd0/ks.cfg**

   You should see messages about formatting the file system and reading the package list. The packages should install, with the only intervention required being to change CDs. Next you should see a post-install message. Finally, you should see the "Complete" message.
4. Remove the floppy; then press the Spacebar to restart your computer (the CD should eject automatically).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip To not have to change CDs with kickstart, do a hard disk or network work install. You can install using kickstart over NFS (ks:nfs:*server*:*path*), from Web server (ks=http://*server*/*path*

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Special Installation Topics

Some things that you run into during installation merit whole discussions by themselves. Rather than bog down the procedures with details that not everyone needs, I have added these topics to this section. Descriptions cover things such as reclaiming disk space and partitioning.

Partitioning your disks

The hard disk (or disks) on your computer provides the permanent storage area for your data files, applications programs, and the operating system (such as Red Hat Linux). Partitioning is the act of dividing a disk into logical areas that can be worked with separately. There are several reasons you may want to do partitioning:

Multiple operating systems — If you install Red Hat Linux on a PC that already has a Windows operating system, you may want to keep both operating systems on the computer. To run efficiently, they must exist on completely separate partitions. When your computer boots, you can choose which system to run.

Multiple partitions within an operating system — To protect from having your entire operating system run out of disk space, people often assign separate partitions to different areas of the Red Hat Linux file system. For example, if `/home` and `/var` were assigned to separate partitions, then a gluttonous user who fills up the `/home` partition wouldn't prevent logging daemons from continuing to write to log files in the `/var/log` directory.

Multiple partitions also make it easier to do certain kinds of backups (such as an image backup). For example, an image backup of `/home` would be much faster (and probably more useful) than an image backup of the root file system (`/`).

Different file system types — Different kinds of file systems that have different structures. File systems of different types must be on their own partitions. In Red Hat Linux, you need at least one file system type for `/` (typically ext3) and one for your swap area. File systems on CD-ROM use the iso9660 file system type.

Tip When you create partitions for Red Hat Linux, you will usually assign the file system type as Linux native (using the ext3 type). Reasons to use other types include needing a file system that allows particularly long filenames or many inodes (each file consumes an inode).

For example, if you set up a news server, it can use many inodes to store news articles. Another reason for using a different file system type is to copy an image backup tape from another operating system to your local disk (such as one from an OS/2 or Minix operating system).

If you have only used Windows operating systems before, you probably had your whole hard disk assigned to C: and never thought about partitions. With Red Hat Linux, you can do a Server class of install (and have Linux erase the whole disk, take it over, and partition it) or a Personal Desktop or Workstation class (and have Linux keep separate partitions for Windows 9*x*/2000 and Linux). With the latest version of Red Hat Linux, you also have the opportunity to view and change the default partitioning for the different installation types.

During installation, Red Hat enables you to partition your hard disk using the Disk Druid utility (a graphical partitioning tool) . The following sections describe how to use Disk Druid (during installation) or fdisk (when Red Hat Linux is up and running). See the section "Tips for creating partitions" for some ideas for creating disk partitions.

Partitioning with Disk Druid during installation

During installation, you are given the opportunity to change how your hard disk is partitioned. Red Hat recommends using the Disk Druid. The Disk Druid screen is divided into two sections. The top shows general information about each hard disk. The bottom shows details of each partition. Figure 2-2 shows an example of the Disk Druid window.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

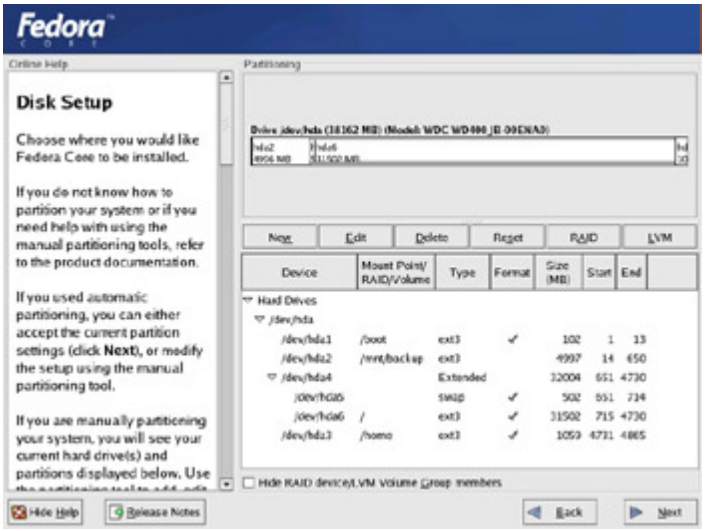

Figure 2-2: Partition your disk during installation from the Disk Setup window.

For each of the hard disk partitions, you can see:

Device — The device name is the name representing the hard disk partition in the `/dev` directory. Each disk partition device begins with two letters: hd for IDE disks, sd for SCSI disks, ed for ESDI disks, or xd for XT disks. After that is a single letter representing the number of the disk (disk 1 is a, disk 2 is b, disk 3 is c, and so on). The partition number for that disk (1, 2, 3, and so on) follows that.

Mount Point/Raid/Volume — The directory where the partition is connected into the Linux file system (if it is). You must assign the root partition (`/`) to a native Linux partition before you can proceed. If you are using RAID or LVM, the name of the RAID device or LVM volume appears here.

Type — The type of file system that is installed on the disk partition. In most cases, the file system will be Linux (ext3), Win VFAT (vfat), or Linux swap. However, you can also use the previous Linux file system (ext2), physical volume (LVM), or software RAID.

Format — Indicates whether (check mark) or not (no check mark) the installation process should format the hard disk partition. Partitions marked with a check are erased! So, on a multiboot system, be sure your Windows partitions, as well as other partitions containing data are not checked!

Size (MB) — The amount of disk space allocated for the partition. If you selected to let the partition grow to fill the existing space, this number may be much larger than the requested amount.

Start/**End** — Represents the partition's starting and ending cylinders on the hard disk.

In the top section, you can see each of the hard disks that are connected to your computer. The drive name is shown first. The Geometry section (Geom) shows the numbers of cylinders, heads, and sectors, respectively, on the disk. That's followed by the model name of the disk. The total amount of disk space, the amount used, and the amount free are shown in megabytes.

Reasons for partitioning

There are different opinions about how to divide up a hard disk. Here are some issues:

Do you want to install another operating system? If you want Windows on your computer along with Linux, you will need at least one Windows (Win95 FAT16 or VFAT type), one Linux (Linux ext3), and one Linux swap partition.

Is it a multiuser system? If you are using the system yourself, you probably don't need many partitions. One reason for partitioning an operating system is to keep the entire system from running out of disk space at once. That also serves to put boundaries on what an individual can use up in his or her home directory (although disk quotas are good for that as well).

Do you have multiple hard disks? You need at least one partition per hard disk. If your system has two hard disks, you may assign one to `/` and one to

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

`/home` (if you have lots of users) or `/var` (if the computer is a server sharing lots of data).

Deleting, adding, and editing partitions

Before you can add a partition, there needs to be some free space available on your hard disk. If all space on your hard disk is currently assigned to one partition (as it often is in DOS or Windows), you must delete or resize that partition before you can claim space on another partition. The section on reclaiming disk space discusses how to add a partition without losing information in your existing single-partition system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution Make sure that any data that you want to keep is backed up before you delete the partition. When you delete a partition, all its data is gone.

The Disk Druid is less flexible, but more intuitive, than the `fdisk` utility. Disk Druid lets you delete, add, and edit partitions.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you create multiple partitions, make sure that there is enough room in the right places to complete the installation. For example, most of the Linux software is installed in the /usr dir

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To delete a partition in Disk Druid, do the following:

1. Select a partition from the list of Current Disk Partitions on the main Disk Druid window (click it or use the arrow keys).
2. To delete the partition, click Delete.
3. When asked to confirm the deletion, click Delete.
4. If you made a mistake, click Reset to return to the partitioning as it was when you started Disk Druid.

To add a partition in Disk Druid, follow these steps from the main Disk Druid window:

1. Select New. A window appears, enabling you to create a new partition.
2. Type the name of the Mount Point (the directory where this partition will connect to the Linux file system). You need at least a root (/) partition and a swap partition.
3. Select the type of file system to be used on the partition. You can select from Linux native (ext2 or preferably ext3), software RAID, Linux swap (swap), physical volume (LVM), or Windows FAT (vfat).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip To create a different file system type than those shown, leave the space you want to use free for now. After installation is complete, use fdisk to create a partition of the type you

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

4. Type the number of megabytes to be used for the partition (in the Size field). If you want this partition to grow to fill the rest of the hard disk, you can put any number in this field (1 will do fine).
5. If you have more than one hard disk, select the disk on which you want to put the partition from the Allowable Drives box.
6. Type the size of the partition (in megabytes) into the Size (MB) box.
7. Select one of the following Additional Size Options:

   Fixed size — Click here to use only the number of megabytes you entered into the Size box when you create the partition.
   Fill all space up to (MB) — If you want to use all remaining space up to a certain number of megabytes, click here and fill in the number. (You may want to do this if you are creating a VFAT partition up to the 2048MB limit that Disk Druid can create.)
   Fill to maximum allowable size — If you want this partition to grow to fill the rest of the disk, click here.
8. Optionally select Force to Be a Primary Partition if you want to be sure to be able to boot the partition or Check for Bad Blocks if you want to have the partition checked for errors.
9. Select OK if everything is correct. (The changes don't take effect until several steps later when you are asked to begin installing the packages.)

To edit a partition in Disk Druid from the main Disk Druid window, follow these steps:

1. Click the partition you want to edit.
2. Click the Edit button. A window appears, ready to let you edit the partition definition.
3. Change any of the attributes (as described in the add partition procedure). For a new install, you may need to add the mount point (/) for your primary Linux partition.
4. Select OK. (The changes don't take effect until several steps later, when you are asked to begin installing the packages.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you want to create a RAID device, you need to first create at least two RAID partitions. Then click the RAID button to make the two partitions into a RAID device. For more information

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Partitioning with fdisk

The `fdisk` utility does the same job as Disk Druid, but it's no longer offered as an option during Red Hat installation. (If you are old school, however, you could press Ctrl+Alt+F2 during the installation process and run `fdisk` from the shell to partition your disk.)

The following procedures are performed from the command line as root user.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution Remember that any partition commands can easily erase your disk or make it inaccessible. Back up critical data before using any tool to change partitions! Then be very careful

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The `fdisk` command is one that is available on many different operating systems (although it looks and behaves differently on each). In Linux, `fdisk` is a menu-based command. To use `fdisk` to list all your partitions, type the following (as root user):

fdisk –l

```
Disk /dev/hda: 40.0 GB, 40020664320 bytes
255 heads, 63 sectors/track, 4865 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot    Start       End    Blocks   Id  System
/dev/hda1   *         1        13    104391   83  Linux
/dev/hda2            14      4833  38716650   83  Linux
/dev/hda3          4834      4865    257040   82  Linux swap
```

To see how each partition is being used on your current system, type the following:

df –h

```
Filesystem            Size  Used Avail Use% Mounted on
/dev/hda2             37G  5.4G   30G  16% /
/dev/hda1             99M  8.6M   86M  10% /boot
none                  61M     0   61M   0% /dev/shm
```

From the output of `df`, you can see that the root of your Linux system (`/`) is on the `/dev/hda2` partition and that the `/dev/hda1` partition is used for `/boot`.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If this had been a dual-boot system (with Windows 98), you might have seen a Windows partition from fdisk that looked like the following:

/dev/hda1 * 1 8 3

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution Before using fdisk to change your partitions, I strongly recommend running the df –h command to see how your partitions are currently being defined. This will help reduce the risk of

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To use `fdisk` to change your partitions, begin (as root user) by typing:

fdisk *device*
where *device* is replaced by the name of the device you want to work with. For example, here are some of your choices:

`/dev/hda`	For the first IDE hard disk; `hdb`, `hdc`, and so on for other IDE disks.
`/dev/sda`	For the first SCSI hard disk; `sdb`, `sdc`, and so on for other SCSI disks.
`/dev/rd/c0d0`	For a RAID device.
`/dev/ida/c0d0`	Also for a RAID device.

After you have started `fdisk`, type **m** to see the options. Here is what you can do with `fdisk`:

Delete a partition — Type **d** and you are asked to enter a partition number on the current hard disk. Type the partition number and press Enter. For example, `/dev/sda2` would be partition number 2. (The deletion won't take effect until you write the change. Until then, it's not too late to back out.)

Create a partition — If you have free space, you can add a new partition. Type **n** and you are asked to enter `l` for a logical partition (5 or over) or `p` for a primary partition (1–4). Enter a partition number from the available range. Then choose the first cylinder number from those available. (The output from `fdisk -l` shown earlier will show you cylinders being used under the Start and End columns.)

Next, enter the cylinder number the partition will end with (or type the specific number of megabytes or kilobytes you want: for example, `+50M` or `+1024K`). You just created an ext3 Linux partition. Again, this change isn't permanent until you write the changes.

Change the partition type — Press `t` to choose the type of file system. Enter the partition number of the partition number you want to change. Type the number representing the file system type you want to use in hexadecimal code. (Type **L** at this point to see a list of file system types and codes.) For a Linux file system, use the number `83`; use `82` for a Linux swap partition. For a windows FAT32 file system, you can use the letter `b`.

Display the partition table — Throughout this process, feel free to type **p** to display (print on the screen) the partition table as it now stands.

Saving and quitting — If you don't like a change you make to your partitions, press **q** to exit without saving. Nothing will have changed on your partition table.

Before you write your changes, display the partition table again and make sure that it is what you want it to be. To write your changes to the partition table, press **w**. You are warned about how dangerous it is to change partitions and asked to confirm the change.

An alternative to the `fdisk` command is `sfdisk`. The `sfdisk` command is command-line–oriented. Type the full command line to list or change partitions. (See the `sfdisk` man page for details.)

Tips for creating partitions

Changing your disk partitions to handle multiple operating systems can be very tricky. Part of the reason is that each different operating system has its own ideas about how partitioning information should be handled, as well as different tools for doing it. Here are some tips to help you get it right.

If you are creating a dual-boot system, particularly for Windows ME or Windows XP, try to install the Windows operating system first. Otherwise, the Windows installation may make the Linux partitions inaccessible.

The `fdisk man` page recommends that you use partitioning tools that come with an operating system to create partitions for that operating system. For example, the DOS `fdisk` knows how to create partitions that DOS will like, and the Red Hat Linux `fdisk` will happily make your Linux partitions. Once your hard disk is set up for dual boot, however, you should probably not go back to Windows-only partitioning tools. Use Linux `fdisk` or a product made for multi-boot systems (such as Partition Magic).

You can have up to 63 partitions on an IDE hard disk. A SCSI hard disk can have up to 15 partitions. You won't need nearly that many partitions.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

If you are using Red Hat Linux as a desktop system, you probably don't need a lot of different partitions within your Red Hat Linux system. There are, however, some very good reasons for having multiple partitions for Red Hat Linux systems that are shared by a lot of users or are public Web servers or file servers. Multiple partitions within Red Hat Linux offer the following advantages:

Protection from attacks — Denial-of-service attacks sometimes take action that tries to fill up your hard disk. If public areas, such as `/var`, are on separate partitions, a successful attack can fill up a partition without shutting down the whole computer. Because `/var` is the default location for Web and FTP servers, and therefore might hold a lot of data, often entire hard disks are assigned to the `/var` file system alone.

Protection from corrupted file systems — If you have only one file system (`/`), corruption of that file system can cause the whole Red Hat Linux system to be damaged. Corruption of a smaller partition can be easier to correct and can often allow the computer to stay in service while the corruption is fixed.

Here are some directories that you may want to consider making into separate file system partitions.

/boot — Sometimes the BIOS in older PCs can access only the first 1024 cylinders of your hard disk. To make sure that the information in your `/boot` directory is accessible to the BIOS, create a separate disk partition (of only a few MB) for `/boot` and make sure that it exists below cylinder 1024. Then, the rest of your Linux system can exist outside of that 1024-cylinder boundary if you like. Even with several boot images, there is no need for `/boot` to be larger than 64MB. For newer hard disks, you can sometimes avoid this problem by selecting "linear mode" during installation. Then the boot partition can be anywhere on the disk.

/usr — This directory structure contains most of the applications and utilities available to Red Hat Linux users. Having `/usr` on a separate partition lets you mount that file system as read-only after the operating system has been installed. This prevents attackers from replacing or removing important system applications with their own versions that may cause security problems. A separate `/usr` partition is also useful if you have diskless workstations on your local network. Using NFS, you can share `/usr` over the network with those workstations.

/var — Your FTP (`/var/ftp`) and Web-server (`/var/www`) directories are, by default, stored under `/var`. Having a separate `/var` partition can prevent an attack on those facilities from corrupting or filling up your entire hard disk.

/home — Because your user account directories are located in this directory, having a separate `/home` account can prevent an indiscriminate user from filling up the entire hard disk. (Disk quotas, see Chapter 10, represent another way of controlling disk use.)

/tmp — Protecting `/tmp` from the rest of the hard disk by placing it on a separate partition can ensure that applications that need to write to temporary files in `/tmp` are able to complete their processing, even if the rest of the disk fills up.

Although people who use Red Hat Linux casually rarely see a need for lots of partitions, those who maintain and occasionally have to recover large systems are thankful when the system they need to fix has several partitions. Multiple partitions can localize deliberate damage (such as denial-of-service attacks), problems from errant users, and accidental file system corruption.

Reclaiming disk space from existing partitions

Chances are that your PC came with a Windows operating system already installed on the entire disk. Installing Red Hat while keeping Windows on your hard disk presents a problem. Red Hat Linux and Windows operating systems are put on separate disk partitions, and right now there is only one partition (and it isn't for Red Hat Linux).

If you are in this predicament, but you have a lot of unused space in your Windows partition (at least 1GB or more), follow this procedure to reclaim the disk space.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Cross-Reference	If, instead of adding Linux to an existing Windows computer, you want to add Windows to your Linux computer, refer to the Linux+Windows+Grub HOWTO. For information on installing Windows NT and Linux on the same computer, refer to the Linux+Windows NT mini-HOWTO.These and other Linux HOWTOs are available from the Linux Documentation Project (www.tldp.org).

Before you begin

Read this section before you begin resizing your Windows partitions!

FIPS is a well-tested utility for resizing DOS FAT and VFAT file systems so that space can be made available to install Linux. To find out what kind of file system is being used on your Windows computer, do the following while Windows is running:

1. Double-click the My Computer icon.
2. Right-click the hard-disk (C:) icon and select Properties.
3. Look at the File system line. If it shows the file system type as FAT or FAT32 the disk partition can be resized with FIPS.

If you are using a Windows 95, Windows 98, or Windows ME computer, providing that you have enough free disk space on the partition, you can probably use FIPS to resize it. If, however, you have a Windows 2000, Windows NT, or Windows XP computer with an NTFS file system, you will have to use some other tools to resize it. Here are some of your choices:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution I have not extensively tested the utilities in the following list. I recommend that before deciding to commit to using any of these tools you research them carefully.

`parted` command — A Red Hat Linux utility called `parted` can be used to resize NTFS file systems. Because it runs in Linux, you need to be able to run it from a boot disk if you currently have a Windows-only system. You can download a boot disk image from the following location and create a boot disk with that

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

image using the boot-disk–creation descriptions in this chapter:
ftp://ftp.gnu.org/gnu/parted/bootdisk/
The README file in this directory describes which boot image to use and how to use it.
Documentation of this utility is available from `/usr/share/doc/parted*/USER`.

Partition Magic — I've heard good reports from people using Partition Magic (`www.partitionmagic.com/partitionmagic`) to resize NTFS partitions. Partition Magic also helps you create new partitions and manage them. It supports Windows XP Professional/Home, Windows 95b-98SE, Windows Me, Windows 2000 Professional, and NT 4.0 workstation (SP6a). The cost is currently $70.

Acronis OS Selector — This is another well-regarded product for managing, creating, and resizing partitions. It supports a variety of file system types, including FAT12, FAT16, FAT32, NTFS, and Linux partition types (ext2, ext3 and Linux ReiserFS). It also supports the same Windows platforms that Partition Magic does. The cost is currently $45 from `www.acronis.com`.

As I mentioned earlier, if your Windows partition is FAT or VFAT you can continue with the procedure in the next session to resize your hard disk.

Backing up and repartitioning

One way to divide a disk that is currently totally used for DOS (DOS, Windows 9*x*/2000, and so on) is not very complicated, but it can take some time. The method is to do a full system backup of your data, erase the whole disk, repartition the disk, restore your Windows 9*x*/2000 operating system to the hard disk (on a smaller partition), and add Red Hat Linux to the hard disk (using new partitions). Whew!

The problem with the backup/repartition approach is that it can be a pain. If your backup device is a 1.44MB floppy disk and your hard disk holds a gigabyte or more, a full backup is not a pleasant prospect. Alternatively, there is a DOS utility called FIPS.

Using the FIPS utility

The FIPS utility came about to solve the problem of a would-be Linux user with a monolithic DOS/Windows hard disk. With FIPS, you can change the size of your DOS partition without erasing it. Though this process is risky, and nobody recommends it without major "buyer beware" warnings, many use it safely and save a lot of time and effort.

FIPS works by changing the values in the partition table and boot sector. Space is gained by changing the partition table that is used to create a new primary DOS partition. After that, the new DOS partition can be converted to Linux partitions.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution Warnings that come with FIPS documentation have mostly to do with older, nonstandard hard disks. I recommend reading the warnings that come with this documentation. You

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Split your DOS (Windows) disk into separate DOS and Linux partitions as follows:

1. **Check your Windows/DOS partition.**
   In a DOS window, run `chkdsk` to check for disk errors. The output shows how your disk space is used, as well as errors found and corrected. Next, run `scandisk` in Windows 9*x* /2000. Scandisk deals with lost file fragments, invalid filenames, and cross-linked files. It also scans for disk surface errors. The point is to correct hard disk errors encountered.
2. **Create a bootable FIPS floppy.**

   You need to create a bootable floppy disk. Insert a blank, 3.5-inch floppy disk in drive A. Then, from DOS (or a DOS window), type the following to create a boot floppy:

   format a:/s
   After creating the boot floppy, copy the following files from the first Red Hat Linux installation CD (CD #1) to this bootable floppy: `restorrb.exe`, `fips.exe`, and `errors.txt`. (These files are located in the `dosutils/fips20` directory.) To make sure that the bootable floppy is working properly, insert the floppy disk into a PC that has a DOS (Windows 9*x* /2000) partition and reboot the computer. Now make sure that you can access the partition by typing **dir C:** at the prompt.
3. **Defragment your hard disk.**
   Reboot your PC so that Windows 9*x*/2000 starts up. To defragment your disk, so that all of your used space is put in order on the disk, open My Computer, right-click your hard disk icon (C:), select Properties, click Tools, and select Defragment Now.

   Defragmenting your disk can be a fairly long process. The result of defragmentation is that all the data on your disk are contiguous, creating a lot of contiguous free space at the end of the partition. There are cases where you will have to do the following special tasks to make this true:

   If the Windows swap file is not moved during defragmentation, you must remove it. Then, after you defragment your disk again and run FIPS, you will need to restore the swap file. To remove the swap file, open the Control Panel, open the System icon, and then click the Performance tab and select Virtual Memory. To disable the swap file, click Disable Virtual Memory.
   If your DOS partition has hidden files that are on the space you are trying to free up, you need to find them. In some cases, you won't be able to delete them. In other cases, such as swap files created by a program, you can safely delete those files. This is a bit tricky because some files should not be deleted, such as DOS system files. You can use the `attrib -s -h` command from the root directory to deal with hidden files.
4. **Reboot (FIPS boot disk).**
   Before you reboot to the FIPS disk, you want to make sure that nothing is written to the hard disk when DOS/Windows shuts down. Look for programs in `config.sys` and `autoexec.bat` that write to disk. When you are satisfied, insert the boot floppy and reboot.
5. **Run FIPS.**

   With your computer now booted in DOS (with C: drive accessible), run the FIPS program by typing:
   # **fips**

   You can quit FIPS at any time by pressing Ctrl+C. If you have more than one hard disk, FIPS will ask which one you want to use. FIPS then displays the partition table and asks you the number of the partition you want to split:

   Type the number of the partition you want to split. FIPS checks the partition for free space and asks if you want to make a copy of your root and boot sector before you go on (this is recommended).
   Type **Y** to backup your root and boot sector. FIPS asks if you have your FIPS bootable floppy in the drive.
   Make sure the floppy is in the drive and type **Y**. FIPS copies a file called `rootboot.000` to your A: drive. Then FIPS determines how much free space is available on the partition. If FIPS can't find at

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

least one cylinder free, it probably means that you need to remove a mirror or image file (after you do that, you will have to defragment the drive again before you can rerun FIPS). If all goes well, you will see output that looks similar to the following:

Old partition	Cylinder	New Partition
2753.3 MB	351	35393.2 MB

Note If the New Partition doesn't show enough disk space, you may need to go back and remove some files until there is enough free space.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Use Right Arrow and Left Arrow keys to choose what cylinder the new partition should start on. As you press those keys, you can see the sizes of the old and new partitions change. When the partitions are the sizes you want, press Enter.
When the new partition table appears, you can accept it or re-edit the table.
Press **C** to continue if everything is okay. You are asked if you are ready to write the new partition table to disk.
Press **Y** to proceed. When FIPS is done, it will say `Bye!` and exit.

6. **Reboot (FIPS boot disk) and test.**
Reboot from the DOS/FIPS disk again. Then run the `fips` command again, but this time in test mode as follows:
fips -t
If the partition table looks okay, exit FIPS by pressing Ctrl+C. You can also use a program such as `chkdsk` to be sure the old partition is still working.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution If the partitions are not correct, you can restore your original partitioning by running the restorrb command from the FIPS floppy disk.

7. **Restart your computer.**
   At this point, you can restart your computer as you would normally (remove the FIPS boot disk). Try using the Windows 9*x*/2000 system to make sure that everything is working okay.

With partitioning done, you can install Red Hat Linux using the reclaimed space. (Don't assign the free disk space to anything yet. Just go ahead and install Red Hat Linux and assign the Linux partitions you need during installation.)
If you have a problem with FIPS, its creator asks that you make a transcript of the FIPS session (using the `-d` option) and send it to `schaefer@rbg.informatik.th-darmstadt.de`. The transcript will appear in the `fipsinfo.dbg` file.

Creating install floppy disks

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

There are several cases in which you need to have one or more floppy disks to install Red Hat Linux. In general, you need an install disk any time the computer you are installing on doesn't have a built-in CD-ROM. This includes network installs and PCMCIA device installs, such as external CD-ROM drives or network adapters on laptops.

Each disk that you may need is available as an image file on the first CD in your Red Hat Linux distribution. If you need to create a disk, copy the image file from the CD to a blank, formatted 1.44MB 3.5-inch floppy disk (using a program that can copy raw data).
The image files are available from the images directory on CD #1 (if the CD were mounted in Red Hat Linux, images are probably in the `/mnt/cdrom/images` directory). Here is a list of the boot images and what each is used for:

Boot (`bootdisk.img`) — Use this boot image as the primary boot image for all install methods.
Block drivers (`drvblock.img`) — Use this disk if the block device you are installing to is not detected by the boot disk.
Network drivers (`drvnet.img`) — Use this disk if the network card on your computer is not detected by the boot disk.
PCMCIA support (`pcmciadd.img`) — Use this image if you need to use a PCMCIA adapter (such as a LAN card on your laptop).

Create the install/boot disks in DOS or Linux. The `rawrite` command for copying the images from DOS to floppy is on Red Hat Linux CD #1. In Linux, use the `dd` command. There are also other image files you can use to create disks to support different types of hardware. These image files include the following:

Block Devices (`drvblock.img`) — Use this image to support special types of hard disks and other block devices.
Network Devices (`drvnet.img`) — Use this image to support different types of network interface cards and other network devices.
PCMCIA Devices (`pcmciadd.img`) — Use this image if you need support for a PCMCIA adapter to install Red Hat Linux that isn't detected by the boot disk.

Creating a floppy disk in Linux

Using the `dd` command, you can copy data from one device to another in Linux (or other UNIX system). To copy a disk image from the installation CD to a disk, you need:

A blank 3.5-inch floppy disk (formatted).
Write permission to the disk drive (probably drive A, which in Red Hat Linux is `/dev/fd0`). The root user should have write permission, but root would have to open write permission to allow another user to write to floppy disk.
An available copy of the disk image to be copied. (The best way to get access to this in Red Hat Linux is by inserting the CD and typing **mount /dev/cdrom**, then typing **cd /mnt/cdrom/images**.)

To copy the `bootdisk.img` file from the images directory on the CD, type the following:

dd if=bootdisk.img of=/dev/fd0 bs=1440k
Replace `bootdisk.img` with the name of any other image file you want to copy.

Creating a disk in DOS

The Red Hat CD has both image files and DOS utilities needed to create the installation disks while you are in DOS or Windows. From a DOS prompt, do the following:

1. Insert Red Hat CD #1 into the drive and change to the CD's drive letter (D: or E:).
2. Insert a blank, formatted 3.5-inch floppy disk into drive A.
3. Change to the `dosutils` directory (`cd \dosutils`).
4. Run the `rawrite` command as follows:
5. C:\> **rawrite**
6. Enter disk image source file name: **D:\images\bootdisk.img**
7. Enter target diskette drive: **a:**
8. Please insert a formatted diskette into drive A: and
   press -Enter- :
9. Press Enter. The image is written to the disk. Remove the disk and mark it.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using GRUB or LILO boot loaders

A boot loader lets you choose when and how to boot the bootable operating systems installed on your computer's hard disks. GRUB is the only boot loader that comes with the Fedora Core distribution of Red Hat Linux. Earlier versions of Red Hat Linux and the current Red Hat Enterprise Linux also include the LILO boot loader. The following sections describe both GRUB and LILO boot loaders.

Booting your computer with GRUB

With multiple operating systems installed and several partitions set up, how does your computer know which operating system to start? To select and manage which partition is booted and how it is booted, you need a boot loader. The boot loader that is installed by default with Red Hat Linux is called the GRand Unified Boot loader (GRUB).
GRUB is a GNU software package (`www.gnu.org/software/grub`) that replaced the LILO boot loader by default in Red Hat Linux 7.2. GRUB offers the following features:

- Support for multiple executable formats.
- Support for Multiboot operating systems (such as Red Hat Linux, FreeBSD, NetBSD, OpenBSD, and other Linux systems).
- Support for non-Multiboot operating systems (such as Windows 95, Windows 98, Windows NT, Windows ME, Windows XP, and OS/2) via a chain-loading function. Chain-loading is the act of loading another boot loader (presumably one that is specific to the proprietary operating system) from GRUB to start the selected operating system.
- Support for multiple file system types.
- Support for automatic decompression of boot images.
- Support for downloading boot images from a network.

For more information on how GRUB works, type **man grub** or **info grub**. The `info` command contains more details about GRUB.

Booting with GRUB

When you install Red Hat Linux, information needed to boot your computer (with one or more operating systems) is automatically set up and ready to go. Simply restart your computer. When you see the GRUB boot screen (it says GRUB at the top and lists bootable partitions below it), do one of the following:

- **Default:** If you do nothing, the default operating system will boot automatically after a few seconds.
- **Select an operating system:** Use the up and down arrow keys to select any of the operating systems shown on the screen. Then press Enter to boot that operating system.
- **Edit the boot process:** If you want to change any of the options used during the boot process, use the arrow keys to select the operating system you want and type **e** to select it. Follow the next procedure to change your boot options temporarily.

If you want to change your boot options so that they take effect every time you boot your computer, see the section on permanently changing boot options. Changing those options involves editing the `/boot/grub/grub.conf` file.

Temporarily changing boot options

From the GRUB boot screen, you can select to change or add boot options for the current boot session. First, select the operating system you want (using the arrow keys) and type **e** (as described earlier). You will see a graphical screen that contains information like the following:

GRUB version 0.93 (639K lower / 128768K upper memory)

root (hd0,0)

kernel /boot/vmlinuz-2.4.22-1 ro root=LABEL=/ hdc=ide-scsi

initrd /boot/initrd-2.4.22-1.img

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Use the (and (keys to select which entry is highlighted.

Press 'b' to boot, 'e' to edit the selected command in the

boot sequence, 'c' for a command-line, 'o' to open a new line

after ('O' for before) the selected line, 'd' to remove the

selected line, or escape to go back to the main menu.

There are three lines in the example of the GRUB editing screen that identify the boot process for the operating system you chose. The first line (beginning with `root`) shows that the entry for the GRUB boot loader is on the first partition of the first hard disk (`hd0,0`). GRUB represents the hard disk as `hd`, regardless of whether it is a SCSI, IDE, or other type of disk. You just count the drive number and partition number, starting from zero.

The second line of the example (beginning with `kernel`) identifies the boot image (`/boot/vmlinuz-2.4.22-1`) and several options. The options identify the partition as initially being loaded `ro` (read-only) and the location of the root file system on a partition with the label LABEL=/. The third line (starting with `initrd`) identifies the location of the initial RAM disk, which contains the minimum files and directories needed during the boot process.

If you are going to change any of the lines related to the boot process, you would probably change only the second line to add or remove boot options. Here is how you do that:

1. Position the cursor on the `kernel` line and type **e**.
2. Either add or remove options after the name of the boot image. You can use a minimal set of bash shell command-line editing features to edit the line. You can even use command completion (type part of a filename and press Tab to complete it). Here are a few options you may want to add or delete:

   Boot to a shell: If you forgot your root password or if your boot process hangs, you can boot directly to a shell by adding `init=/bin/sh` to the boot line. (The file system is mounted read-only, so you can copy files out. You need to remount the file system with read/write permission to be able to change files.)

   Select a run level: If you want to boot to a particular run level, you can add the word *linux,* followed by the number of the run level you want. For example, to have Red Hat Linux boot to run level 3 (multiuser plus networking mode), add `linux 3` to the end of the boot line. You can also boot to single-user mode (1), multi-user mode (2), or X GUI mode (5). Level 3 is a good choice if your GUI is temporarily broken.
3. Press Enter to return to the editing screen.
4. Type **b** to boot the computer with the new options. The next time you boot your computer, the new options will not be saved. To add options so they are saved permanently, see the next section.

Permanently changing boot options

You can change the options that take effect each time you boot your computer by changing the GRUB configuration file. In Red Hat Linux, GRUB configuration centers around the `/boot/grub/grub.conf` file.

The `/boot/grub/grub.conf` file is created when you install Red Hat Linux. Here is an example of a `grub.conf` file.

grub.conf generated by anaconda

#

Note that you do not have to rerun grub after making

changes to this file

NOTICE: You have a /boot partition. This means that

all kernel and initrd paths are relative to /boot/, eg.

root (hd0,0)

kernel /vmlinuz-version ro root=/dev/hda3

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
#       initrd /initrd-version.img
#boot=/dev/hda
default=0
timeout=10
splashimage=(hd0,4)/grub/splash.xpm.gz
title Red Hat Linux (2.4.22-1)
    root (hd0,4)
    kernel /vmlinuz-2.4.22-1 ro root=LABEL=/
    initrd /initrd-2.4.22-1.img
title Windows XP
    rootnoverify (hd0,0)
    chainloader +1
```

The `default=0` line indicates that the first partition in this list (in this case Red Hat Linux) will be the one that is booted by default. The line `timeout=10` causes GRUB to pause for ten seconds before booting the default partition. (That's how much time you have to press `e` if you want to edit the boot line, or to press arrow keys to select a different operating system to boot.)

The `splashimage` line looks in the fifth partition on the first disk (`hd0,4`) for the boot partition (in this case `/dev/hda5`, which is the `/boot` partition). GRUB loads `splash.xpm.gz` as the image on the splash screen (`/boot/grub/splash.xpm.gz`). The splash screen appears as the background of the boot screen.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note GRUB indicates disk partitions using the following notation: (hd0,0). The first number represents the disk, and the second is the partition on that disk. So, (hd0,1) is the second par

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The two bootable partitions in this example are `Red Hat Linux` and `Windows XP`. The title lines for each of those partitions are followed by the name that appears on the boot screen to represent each partition.
For the Red Hat Linux system, the `root` line indicates the location of the boot partition as the second partition on the first disk. So, to find the bootable `kernel` (`vmlinuz-2.4.22-1`) and the `initrd` initial RAM disk boot image that is loaded (`initrd-2.4.22-1.img`), GRUB looks in the root of `hd0,4` (which is represented by `/dev/hda5` and is eventually mounted as `/boot`). Other options on the `kernel` line set the partition as read-only initially (`ro`) and set the root file system to `/dev/hda6`.
For the Windows XP partition, the `rootnoverify` line indicates that GRUB should not try to mount the partition. In this case, Windows ME is on the first partition of the first hard disk `(hd0,0)` or `/dev/hda1`. Instead of mounting the partition and passing options to the new operating system, the `chainloader +1` indicates to hand control the booting of the operating system to another boot loader. The `+1` indicates that the first sector of the partition is used as the boot loader.

Note Microsoft operating systems require that you use the chainloader to boot them from GRUB.

If you make any changes to the `/boot/grub/grub.conf` file, you *do not* need to load those changes. Those changes are automatically picked up by GRUB when you reboot your computer. If you are accustomed to using the LILO boot loader, this may confuse you at first, as LILO requires you to rerun the `lilo` command for the changes to take effect.

Adding a new GRUB boot image

You may have different boot images for kernels that include different features. Here is the procedure for modifying the `grub.conf` file:

1. Copy the new image from the directory in which it was created (such as `/usr/src/linux-2.4/arch/i386/boot`) to the `/boot` directory. Name the file something that reflects its contents, such as `bz-2.4.21`. For example:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

2. # **cp /usr/src/linux-2.4/arch/i386/boot/bzImage /boot/bz-2.4.22**
3. Add several lines to the `/boot/grub/grub.conf` file so that the image can be started at boot time if it is selected. For example:
4. title Red Hat Linux (IPV6 build)
5. root (hd0,4)
6. kernel /bz-2.4.22 ro root=/dev/hda6 hdc=ide-scsi

   initrd /initrd-2.4.22.img
7. Reboot your computer.
8. When the GRUB boot screen appears, move your cursor to the title representing the new kernel and press Enter.

The advantage to this approach, as opposed to copying the new boot image over the old one, is that if the kernel fails to boot, you can always go back and restart the old kernel. When you feel confident that the new kernel is working properly, you can use it to replace the old kernel or perhaps just make the new kernel the default boot definition.

Booting your computer with LILO

LILO stands for LInux LOader. Like other boot loaders, LILO is a program that can stand outside the operating systems installed on the computer so you can choose which system to boot. It also lets you give special options that modify how the operating system is booted.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note LILO is not part of the Fedora Core version of Red Hat Linux, so it is not available on the CDs that come with this book. However, LILO is still available from Red Hat FTP sites with hold

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

If LILO is being used on your computer, it is installed in either the master boot record or the first sector of the root partition. The master boot record is read directly by the computer's BIOS. In general, if LILO is the only loader on your computer, install it in the master boot record. If there is another boot loader already in the master boot record, put LILO in the root partition.

Using LILO

When your computer boots with the Red Hat Linux version of LILO installed in the master boot record, a graphical Red Hat screen appears, displaying the bootable partitions on the computer. Use the Up Arrow and Down Arrow keys on your keyboard to select the one you want and press Enter. Otherwise, the default partition that you set at installation will boot after a few seconds.

If you want to add any special options when you boot, press Ctrl+X. You will see a text-based boot prompt that appears as follows:

boot:

LILO pauses for a few seconds and then automatically boots the first image from the default bootable partition. To see the bootable partitions again, quickly press Tab. You may see something similar to the following:

LILO boot:

linux linux-up dos

boot:
This example shows that three bootable partitions are on your computer, called `linux`, `linux-up`, and `dos`. The first two refer to two different boot images that can boot the Linux partition. The third refers to a bootable DOS partition (presumably containing a Windows operating system). The first bootable partition is loaded if you don't type anything after a few seconds. Or you could type the name of the other partition to have that boot instead.
If you have multiple boot images, press Shift, and LILO will ask you which image you want to boot. Available boot images and other options are defined in the `/etc/lilo.conf` file.

Setting up the /etc/lilo.conf file

The `/etc/lilo.conf` file is where LILO gets the information it needs to find and start bootable partitions and images. By adding options to the `/etc/lilo.conf` file, you can change the behavior of the boot process. The following is an example of some of the contents of the `/etc/lilo.conf` file:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note

Because LILO is not used by default, there is no /etc/lilo.conf file. However, the Red Hat Linux installation program creates an /etc/lilo.conf.anaconda file to suit your installation. If you ch

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
prompt
timeout=50
default=linux
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
message=/boot/message
linear

image=/boot/vmlinuz-2.4.22-1
    label=linux
    initrd=/boot/initrd-2.4.22-1.img
    read-only
    root=/dev/hda6
    append="root=LABEL=/"

other=/dev/hda1
    optional
    label=dos
```

With prompt on, the boot prompt appears when the system is booted without requiring that any keys are pressed. The timeout value, in this case 50 tenths of a second (5 seconds), defines how long to wait for keyboard input before booting the default boot image. The boot line indicates that the bootable partition is on the hard disk represented by `/dev/hda` (the first IDE hard disk).

The `map` line indicates the location of the map file (`/boot/map`, by default). The map file contains the name and locations of bootable kernel images. The `install` line indicates that the `/boot/boot.b` file is used as the new boot sector. The `message` line tells LILO to display the contents of the `/boot/message` file when booting (that contains the graphical Red Hat boot screen that appears). The linear line causes linear sector addresses to be generated (instead of sector/head/cylinder addresses).

In the sample file, there are two bootable partitions. The first (`image=/boot/vmlinuz-2.4.22-1`) shows an image labeled `linux`. The root file system (`/`) for that image is on partition `/dev/hda6`. Read-only indicates that the file system is first mounted read-only, though it is probably mounted as read/write after a file system check. The `inidrd` line indicates the location of the initial RAM disk image used to start the system.

The second bootable partition, which is indicated by the word *other* in this example, is on the `/dev/hda1` partition. Because it is a Windows XP system, it is labeled a DOS file system. The table line indicates the device that contains the partition.

Other bootable images are listed in this file, and you can add another boot image yourself (like one you create from reconfiguring your kernel as discussed in the next section) by installing the new image and changing `lilo.conf`.

After you change `lilo.conf`, you then must run the `lilo` command for the changes to take effect. You may have different boot images for kernels that include different features. Here is the procedure for modifying the `lilo.conf` file:

1. Copy the new image from the directory in which it was created (such as `/usr/src/linux-2.4/arch/i386/boot`) to the `/boot` directory. Name the file something that reflects its contents, such as `zImage-2.4.22-1`.
2. Add several lines to the `/etc/lilo.conf` file so that the image can be started at boot time if it is selected. For example:
3.
```
image=/boot/zImage-2.4.22-1
    label=new
```
4. Type the `lilo -t` command (as root user) to test that the changes were okay.
5. Type the `lilo` command (with no options) for the changes to be installed.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To boot from this new image, either select new from the graphical boot screen or type `new` and press Enter at the LILO boot prompt. If five seconds is too quick, increase the timeout value (such as 100 for 10 seconds).
Options that you can use in the `/etc/lilo.conf` file are divided into global options, per-image options, and kernel options. There is a lot of documentation available for LILO. For more details on any of the options described here or for other options, you can see the `lilo.conf` manual page (type `man lilo.conf`) or any of the documents in `/usr/share/doc/lilo*/doc`.
A few examples follow of global options that you can add to `/etc/lilo.conf`. Global options apply to LILO as a whole, instead of just to a particular boot image.
You can use the `default=label` option, where label is replaced by an image's label name, to indicate that a particular image be used as the default boot image. If that option is excluded, the first image listed in the `/etc/lilo.conf` file is used as the default. For example, to start the image labeled new by default, add the following line to `lilo.conf`:

```
default=new
```

Change the delay from 5 seconds to something greater if you want LILO to wait longer before starting the default image. This gives you more time to boot a different image. To change the value from 5 seconds (50) to 15 seconds (150), add the following line:

```
delay=150
```

You can change the message that appears before the LILO prompt by adding that message to a file and changing the message line. For example, you could create a `/boot/boot.message` file and add the following words to that file: `Choose linux, new, or dos`. To have that message appear before the boot prompt, add the following line to `/etc/lilo.conf`:

```
message=/boot/boot.message
```

All per-image options begin with either an `image=` line (indicating a Linux kernel) or `other=` (indicating some other kind of operating system, such as Windows XP). The per-image options apply to particular boot images rather than to all images (as global options do). Along with the image or other line is a `label=` line, which gives a name to that image. The name is what you would select at boot time to boot that image. Here are some of the options that you can add to each of those image definitions:

lock — This enables automatic recording of `boot` command lines as the defaults for different boot options.
alias=name — You can replace *name* with any name. That name becomes an alias for the image name defined in the label option.
password=password — You can password-protect the image by adding a password option line and replacing *password* your own password. The password would have to be entered to boot the image.
restricted — This option is used with the password option. It indicates that a password should be used only if command-line options are given when trying to boot the image.

For Linux kernel images, there are specific options that you can use. These options let you deal with hardware issues that can't be autodetected, or provide information such as how the root file system is mounted. Here are some of kernel image–specific options:

append — Add a string of letters and numbers to this option that need to be passed to the kernel. In particular, these can be parameters that need to be passed to better define the hard disk when some aspect of that disk can't be autodetected.
ramdisk — Add the size of the RAM disk that you want to use so as to override the size of the RAM disk built into the kernel.
read-only — Indicates to mount the root file system read-only. It is typically remounted read-write after the disk is checked.
read-write — Indicates to mount the root file system read/write.

Changing your boot loader

If you don't want to use the GRUB boot loader, or if you tried out LILO and want to switch back to GRUB, it's not hard to change to a different boot loader. The one issue is that if you need to obtain the lilo package, you're on your own since it's not on the CDs that come with this book. To switch your boot loader from GRUB to LILO, do the following:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

1. Configure the `/etc/lilo.conf` file as described in the "Booting your computer with LILO" section. (Use the contents of `/etc/lilo.conf.anaconda` to start.)
2. As root user from a Terminal window, type the following:
3. # **lilo**
4. The new Master Boot Record is written, including the entries in `/etc/lilo.conf`.
5. Reboot your computer. You should see the LILO boot screen.

To change your boot loader from LILO to GRUB, do the following:

1. Configure the `/boot/grub/grub.conf` file as described in the "Booting your computer with GRUB" section.
2. You need to know the device on which you want to install GRUB. For example, to install GRUB on the master boot record of the first disk, type the following as root user from a Terminal window:
3. # **grub-install /dev/hda**

   The new Master Boot Record is written to boot with the GRUB boot loader.
4. Reboot your computer. You should see the GRUB boot screen.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Troubleshooting Your Installation

After you have finished installing Red Hat Linux, you can check how the installation went by checking your log files. There are three places to look once the system comes up:

`/tmp/upgrade.log` — When upgrading packages, output from each installed package is sent to this file. You can see what packages were installed and if any failed.

`/var/log/dmesg` — This file contains the messages that are sent to the console terminal as the system boots up, including messages relating to the kernel being started and hardware being recognized. If a piece of hardware isn't working, you can check here to make sure that the kernel found the hardware and configured it properly.

`/var/log/boot.log` — This file contains information about each service that is started up at boot time. You can see if each service started successfully. If a service fails to start properly, there may be clues in this file that will help you learn what went wrong.

If something was set wrong (such as your mouse) or just isn't working quite right (such as your video display), you can always go back after Red Hat Linux is running and correct the problem. Here is a list of utilities you can use to reconfigure different features that were set during installation:

Changing or adding a mouse: `redhat-config-mouse`
Changing a keyboard language: `redhat-config-keyboard`
Adding or deleting software packages: `redhat-config-packages` or `rpm`
Partitioning: `fdisk` or `sfdisk`
Boot loader: `/boot/grub/grub.conf` (for GRUB); `lilo` and `/etc/lilo.conf` (for LILO)
Networking (Ethernet & TCP/IP): `redhat-config-network`
Time zone: `timeconfig` or `firstboot`
User accounts: `useradd` or `redhat-config-users`
Sound: `redhat-config-soundcard`
X Window System: `redhat-config-xfree86`

Here are a few other random tips that can help you during installation:

If installation fails because the installation procedure is unable to detect your video card, try restarting installation in text mode. After Red Hat Linux is installed and running, use the `redhat-config-xfree86` command to configure your video card and monitor. (For some cards, such as those from NVIDIA, you need to get and install special drivers from the manufacturer's Web site.)

If installation completes successfully, but your screen is garbled when you reboot, you should try to get Red Hat Linux to boot to a text-login prompt. To do this, add the words `linux 3` to the end of the kernel boot line in GRUB or LILO. Linux will start with the GUI temporarily disabled. Run `redhat-config-xfree86` to try to fix the problem. (See Chapter 3 for other advice related to fixing your GUI.)

If your mouse is not detected during installation, you can use arrow keys and the Tab key to make selections.

Probably the best resource for troubleshooting your installation problems is the Red Hat Support site (`www.redhat.com/apps/support`). Links from that page can take you to documentation, updates and errata, and information about support programs. If you are having problems with a particular piece of hardware, try searching the Solutions Database, using the name of the hardware in the search box. If you are having problems with particular hardware, chances are someone else did, too.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

Installing Linux is not nearly the adventure it once was. Precompiled binary software and preselected packaging and partitions make most Red Hat Linux installations a simple proposition. Improved installation and GUI configuration windows have made it easier for computer users who are not programmers to enter the Linux arena.

Besides providing some step-by-step installation procedures, this chapter discussed some of the trickier aspects of Red Hat Linux installation. In particular, ways of partitioning your hard disk, creating installation boot disks, and changing the boot procedure were described.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Part II: Using Red Hat Linux

Chapter List

Chapter 3: Getting Started with the Desktop
Chapter 4: Using Linux Commands
Chapter 5: Accessing and Running Applications
Chapter 6: Publishing with Red Hat Linux
Chapter 7: Playing Games with Red Hat Linux
Chapter 8: Multimedia in Red Hat Linux
Chapter 9: Tools for Using the Internet and the Web

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 3: Getting Started with the Desktop

Overview

The desktop is the most personal feature of your computer. The way that icons, menus, panels, and backgrounds are arranged and displayed should make it both easy and pleasant to do your work. With Red Hat Linux, you have an incredible amount of control over how your desktop behaves and how your desktop is arranged.

In the past few versions of Red Hat Linux, the desktop has changed dramatically. While going for a "Red Hat" look-and-feel, Red Hat, Inc. set out to standardize how you use Red Hat Linux, regardless of whether you choose KDE or GNOME as your desktop environment. Of course, those who like the old KDE or GNOME may be a bit upset about the change.

Not to be outdone, KDE and GNOME themselves, with recent updates to release 3.1 and 2.4, respectively, have stretched beyond the bounds of a simple look-and-feel. With each desktop environment, you can get a full set of desktop applications, features for launching applications, and tools for configuring preferences.

This chapter takes you on a tour of your desktop — going through the process of logging in, trying out some features, and customizing how your desktop looks and behaves. Sections on KDE and GNOME desktops contain reference information on how to set preferences, run applications, configure panels, and work with the file manager. The last section describes how to use the Display Settings window to configure your video card and monitor if they were not properly detected.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Logging in to Red Hat Linux

Because Red Hat Linux was created as a multiuser computer system, you start by logging in (even if you are the only person using the computer). Logging in accomplishes three functions:

- Identifies you as a particular user.
- Starts up your own shell and desktop (icons, panels, backgrounds, and so on) configurations.
- Gives you appropriate permissions to change files and run programs.

After the computer has been turned on and the operating system has started, you see either a graphical login screen (default) or a text-based login prompt. The text-based prompt should look like this:

Fedora Core release 1

Kernel 2.4.22-1 on an i686

localhost login:

Note If you see a text-based login prompt instead of the graphical login screen, and you want to use the GUI, type your user name and password. Then when you see a command prompt, type **startx** to start up your desktop.

The graphical login is typically your entry into the X Window System graphical user interface (GUI). Figure 3-1 is an example of the login window you see when you boot the Red Hat Linux software that comes with this book (called the Fedora Core).

Figure 3-1: A graphical login can access the Red Hat Linux GUI.

Notice the several menu buttons on the login screen. You can ignore them and simply type your user name and password as prompted. Or you can use these buttons as follows:

- **Language** — Click Option to select a language other than the last language you selected when Red Hat Linux was installed. (You may need additional software packages to use different languages.)
- **Session** — When you log in, your graphical desktop usually starts up (either GNOME or KDE). Click Session to choose a different desktop (if available) or to select Failsafe. If you can't remember which desktop to use, you can choose Last (to use the desktop you used most recently) or Default (to use the desktop set as your default). You use Failsafe if you want only a shell interface (the shell is described in Chapter 4).

Note The Failsafe session is used primarily to correct problems when, for some reason, your desktop session won't start properly. For example, graphics settings may be wrong, resulting in a garbled screen. The Failsafe session simply opens a shell window so you can type commands to correct the problem before you log in again. When you are done with a Failsafe session, type **exit** to be allowed to log in again.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Reboot — Click to reboot the computer.
Shutdown — Click to shut down the computer.

To log in, type your user name and, when prompted, your password. You should log in as either a regular user or as the root user:

A regular user — As someone just using the Linux system, you probably have your own unique user name and password. Often, that name is associated with your real name (such as johnb, susanp, or djones). If you are still not sure why you need a user login, see the sidebar "Why Do I Need a User Login?"

The root user — Every Linux system has a *root* user assigned when Linux is installed. The root user (literally type the user name ***root***) can run programs, use files, and change the computer setup in any way. Because the root user has special powers, and can therefore do special damage, you usually log in as a regular user (which allows access only to that user's files and those that are open to everyone).

Cross-Reference See Chapter 10 for a description of the root user and Chapter 11 for information on how to set up and use other user accounts. Refer to Chapter 14 for suggestions on how to choose a good password.

If your desktop did not start, refer to the "Troubleshooting Your Desktop" section at the end of this chapter. Otherwise, continue on to "Getting Familiar with the Desktop."

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Getting Familiar with the Desktop

Desktop refers to the presentation of windows, menus, panels, icons, and other graphical elements on your computer screen. Originally, computer systems such as Linux operated purely in text mode — no mouse, no colors, just commands typed on the screen. Desktops provide a more intuitive way of using your computer.

[redacted]

Why Do I Need a User Login?

If you [redacted]on a PC, and you are the only one using your Linux computer, you may wonder why you need a user account and password. Unlike Windows, Linux (as its predecessor UNIX) was designed from the ground up to be a multiuser system. Here are several good reasons why you should use separate user accounts:

- Even as the only person using Linux, you want a user name besides root for running applications and working with files so you don't change critical system files by mistake during everyday computer use.
- If several people are using a Linux system, separate user accounts let you protect your files from being accessed or changed by others.
- Networking is probably the best reason for using a Linux system. If you are on a network, a unique user name is useful in many ways. Your user name can be associated with resources on other computers: file systems, application programs, and mailboxes to name a few. Sometimes a root user is not allowed to share resources on remote Linux systems.
- Over time, you will probably change personal configuration information associated with your account. For example, you may add aliases, create your own utility programs, or set properties for the applications you use. By gathering this information in one place, it's easy to move your account or add a new account to another computer in the future.

[redacted]

Like most th[redacted], the desktop is built from a set of interchangeable building blocks. The building blo[redacted]esktop, to use a car analogy, are:

- The X Window System (which is like the frame of the car),
- The KDE or GNOME desktop environment (which is like a blueprint of how the working parts fit together)
- The Metacity window manager (which provides the steering wheel, seat upholstery, and fuzzy dice on the mirror).
- The Bluecurve desktop theme (the paint job and the pin stripe).

Once Red Hat Linux is installed (see Chapter 2) and you have logged in (see the previous section), you should see either the GNOME or KDE desktop with the Red Hat look-and-feel. At this point, I'll take you on a tour of the desktop and step you through some initial setup to get your desktop going. Figure 3-2 shows an example of the Red Hat Linux default desktop (GNOME).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

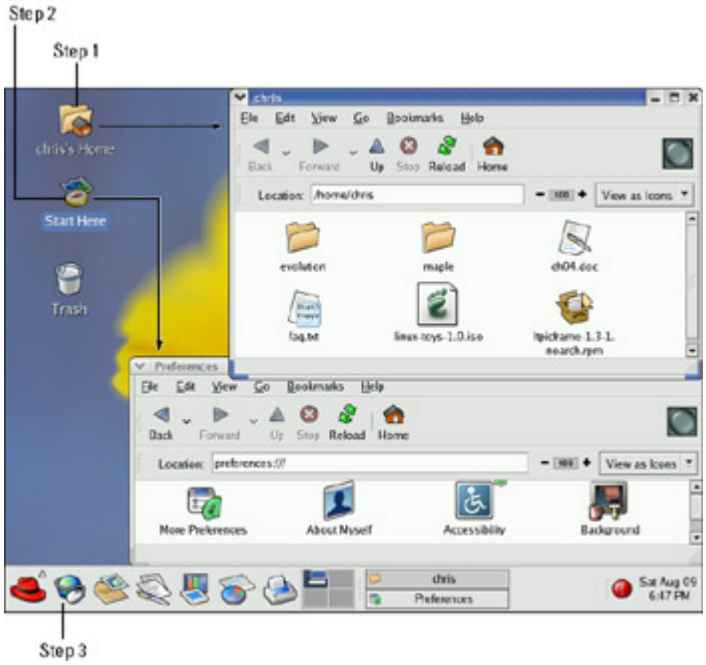

Figure 3-2: After login, Red Hat Linux starts you off with a GNOME desktop by default. Steps 1, 2, and 3 are explained in the following section.

Because GNOME is the default desktop for Red Hat personal desktop and workstation installs, I'll start by walking you around the GNOME desktop. As indicated in Figure 3-2, the tour steps you through trying out your home folder, changing some preferences, and configuring your panel.

To understand how Red Hat has created a common look-and-feel for the GNOME and KDE environments, refer to the sidebar, "The Red Hat Look-and-Feel."

Touring your desktop

If you are unfamiliar with the Red Hat desktop, I suggest you take this quick tour to familiarize yourself with its features:

Step 1: Checking out your home folder

Double-click the "*user*'s Home" icon on the desktop (it should say something like "bill's Home" or "julie's Home," depending on your user name). The window that appears shows your file manager window (Nautilus in GNOME or Konqueror in KDE) as it displays the contents of your home folder. The location of the home folder (also referred to as a *home directory*) on your computer is usually `/home/user`, where *user* is replaced by your user name. Here are some things to try out with your home folder:

The Red Hat Look-and-Feel

Red Hat has done a lot in the past few versions to make KDE and GNOME appear to be the same in many ways. Despite the fact that the differences between KDE and GNOME below the surface are quite dramatic, the way icons, menus, panels, and many Red Hat system tools are presented are the same for the two environments.

To provide a common interface in two desktop environments, you will find the following similarities between the KDE and GNOME desktops in Red Hat Linux 10:

Main menu — An icon on the far left of the panel (in both KDE and GNOME) represents the main menu used in both desktop environments. The menus in both desktops contain similar content, although GNOME represents this menu with a Red Hat icon and KDE uses a K icon.

Panel arrangement — On their panels, both KDE and GNOME start with the same main menu button (just described), the same six desktop applications, a viewer for

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

virtual windows, an area for listed tasks, and a clock.

System tools, system settings, and server settings — These three menus listed on the main menu contain most of the same applications for both desktop environments. This is where Red Hat has added many tools from its growing list of administrative GUIs.

Icons — Both KDE and GNOME now start with at least Home, StartHere, and Trash icons on the desktop. These desktop icons, as well as icons shown on the panel and menus, are Red Hat Linux–specific. The Red Hat icons add a lot to the similar look of both KDE and GNOME in Red Hat Linux.

Although the two desktops now look very much alike in Red Hat Linux and offer many of the same Red Hat administrative tools, the two desktop environments are very different in other ways. KDE comes with many more integrated applications and offers more tools for configuring desktop preferences. GNOME tends to be more streamlined, offering a fairly simple, efficient desktop experience.

To choose between KDE and GNOME, I recommend you read the sections on the two desktop environments and compare features in their panels, file managers (Konqueror and Nautilus), and configurable preferences.

1. **Folders** — Create folders and subfolders to store your work (click File → New Folder, and then name the new folder something like Images, Music, or Packages).
2. **Open With** — Click the right mouse button on any object in a folder, then select Open With. You should be able to see several programs you can use to open the object. For example, you can choose to open a Web page (.html file) with htmlview, Lynx, or Mozilla (provided each of those tools is installed).
3. **Side Pane** — Click View → Side Pane. From the box at the top of this pane, choose Information to show information about the selected folder or file. Next choose History, to see files and folders previously viewed.
4. **Backgrounds** — Click Edit → Backgrounds and Emblems. Drag-and-drop patterns or colors you like into the pane on your folder window. (I personally like camouflage.) Click Emblems, then drag-and drop an emblem on a file or folder. Use the emblem to remind yourself of something about the object (such as the fact that it's a personal document or is of an urgent nature).
5. **Organize your work** — As you create documents, add music, or download images from your camera, organize them into your home folder or any subfolders. Your home folder is not accessible to any other user on the computer except the root user, so you can safely store your work there. With the files you create, you can:

   Move — Drag-and-drop to move a file to another folder icon or folder window.

   Delete — Drag-and-drop a file to the Trash icon to delete it.

   Rename — Right-click on the file, select Rename, and then type the new name.

As with any window, with the Folder window you can:

Window shade — Double-click in the title bar to roll the window up and down.

Minimize/Unminimize — Click the Minimize button (first button, upper-right corner of the title bar) to minimize the window to the window pane. Click the minimized window in the desktop panel to return it to your desktop.

Maximize — Click the Maximize button (second button, upper-right corner of the title bar) to have the window go full screen.

Delete — Click the X button (upper-right corner of the title bar) to delete the window.

Step 2: Change some preferences

More than 20 different preference categories are available from the GNOME desktop. Double-click the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Start Here icon, and then double-click the Preferences icon. Here are a few preferences you might want to modify when you start out (see the GNOME and KDE preferences section later in this chapter for further details):

Change background — Double-click the Background icon. The Background Preferences window appears as shown in Figure 3-3.To change the background image, select one of the Picture Options. To change the image, click the Select Picture box (to choose a file from your disk) or drag-and-drop an image onto that box. For example, from a folder window, type **`/usr/share/backgrounds/images`** in the Location box. Then drag-and-drop images to the Select Picture box to change the background.

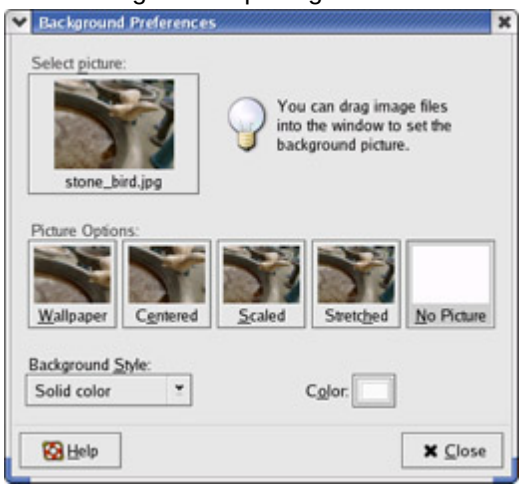

Figure 3-3: Set a colors or pictures for your desktop background.

Choose browsers and editors — Double-click the Preferred Applications icon. When you open a Web page, text file, or shell from the desktop, this preferences window lets you choose which Web browser, text editor, or Terminal window to open. Mozilla is the default, but you can choose Konqueror (the KDE browser) to run on your GUI. If you want a text-based Web browser, select Lynx or Links. To use a different Web browser, such as epiphany, select Custom and type the name of that browser. For your default text editor, choose the application to open plain-text files from your desktop (by default, OpenOffice is used). The GNOME Terminal is used when you need a shell prompt (you can change that to use a standard xterm Terminal).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip The konsole command starts the Konsole (KDE) terminal window. Programmers who use many terminal windows at once often prefer Konsole over gnome-terminal (finding it to be

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Add a screensaver — Double-click the Screensaver icon. Try out a few screen savers (all are used randomly, by default). Click "Only One Screen Saver" in the Mode box.

Click different screen savers to see them and click Preview to try them out. Try T3D for a bizarro clock or Rubick to watch how to solve a Rubick's cube.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Click the Lock Screen After check box and set the number of minutes after which the screen will be locked. This is a good option for an office environment, where you want your scr

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Change the theme — Double-click the Theme icon. You can change the entire theme (colors, icons, borders, and so on) for your desktop. The Bluecurve theme is the default theme used by Red Hat Linux. Try any of the others to find one that suits you. Click Theme Details to mix-and-match attributes from different themes. Figure 3-4 shows the window for selecting a theme.

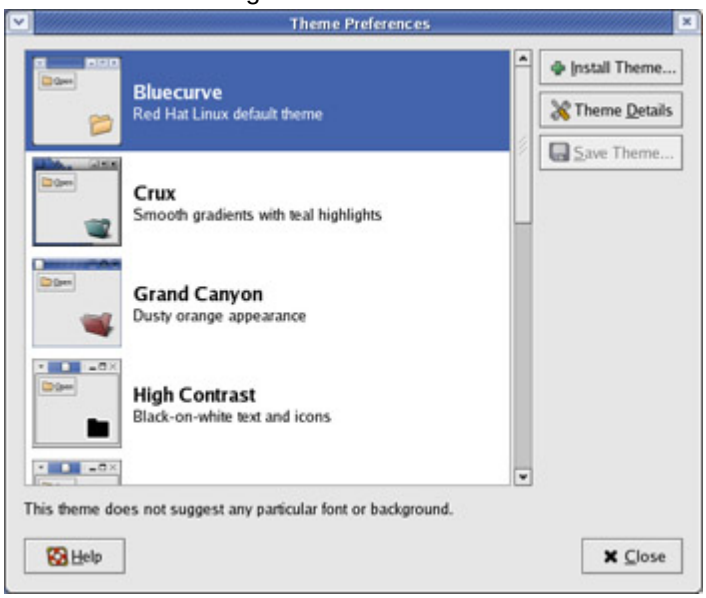

Figure 3-4: Change from the default Bluecurve theme.

Step 3: Configure your panel

Most people manage their desktops from the panel that appears at the bottom of the screen. This panel provides a natural way to:

Launch applications
Change workspaces
Add useful information (clocks, news tickers, CD players, and so on)

Step through the following procedure to learn about the Red Hat (GNOME) panel:

1. **Red Hat Menu** — Click the red hat in the panel. Most useful GUI applications and system tools that come with Red Hat Linux are available from the menus and submenus of this main Red Hat menu. Here are a few cool things to do from this menu:
   Start an application — Click Accessories, Games, Graphics, Internet, Office, or Sound & Video menu items, then select any application to run.
   Change your settings — Click the Preferences menu item to change preferences or the System Settings or System Tools menu items to change system-wide settings.
   Log out or shutdown — Click the Logout menu item to log out from your current desktop session, shut down, or reboot your computer.
2. **Launch applications** — Red Hat places icons for popular desktop applications right on the panel. Click any of the icons shown in Figure 3-5 to launch a Web browser (Mozilla), e-mail reader (Evolution), word processor (OpenOffice Writer), presentation creator (OpenOffice Impress), spreadsheet (OpenOffice Calc), or print manager, respectively.

Figure 3-5: Launch popular desktop applications with one click.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

3. **Use workspaces** — Click different panels in the Workspace Switcher. Open an application, then click on another workspace panel. Workspaces are a great way to have multiple windows and still keep your desktop uncluttered. Notice that there are tiny representations of each window you open on the workspace panel it is in. Drag and drop the tiny windows to move them to different workspaces, without leaving your workspace.
4. **Add cool stuff to your panel** — Click on an empty place in the panel so that a panel menu appears. It should say "Add to Panel" at the top. Because real estate is limited on your panel, I recommend adding a drawer, to which you can add some little applications that run in the panel and icons that launch other applications. To begin, click Add to Panel → Drawer. A Drawer icon appears on your panel (you can drag it where you want it to go. Right-click the drawer, then click Add to Panel. Here are a few things I suggest adding to your panel:

    Terminal — From the drawer menu, click Add to Panel →Launcher from menu → System Tools → Terminal. Now, when I ask you to type something into a Terminal window, you can launch one from this drawer.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Throughout this book, I give examples that require you to use a Terminal window. Neither the new KDE nor GNOME desktops have a Terminal window launcher on the panel or

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Weather report — From the drawer menu, click Add to Panel → Accessories → Weather Report. Right-click on the temperature icon that appears and choose Preferences. Click the Location tab and select your country, state, and city from the list. Now, whenever you double-click the temperature icon in your drawer, you can see weather conditions and a forecast for your city.
Volume control — From the drawer menu, click Add to Panel → Multimedia → Volume Control. The volume control that appears lets you control audio levels to your sound card (when you play a CD, for example).
Popular folder or Web site — Folders or Web sites that you visit often should be easily accessible. Click Add to Panel → Launcher. From the Create Launcher window, select Link (as the Type) and type a URL or Application, and then type **nautilus *folder*** (where *`folder`* is replaced by the folder you want to open). Click Icon and choose an icon to represent the item. For example, when I write a book, I have a folder containing chapter files. To add a launcher for that file, I select Application from Launcher Properties (`nautilus /home/chris/rhlbible`). Then I assign an icon (an image of the book cover) to the launcher. Figure 3-6 shows an example of a drawer, with the launchers I just described added to it.

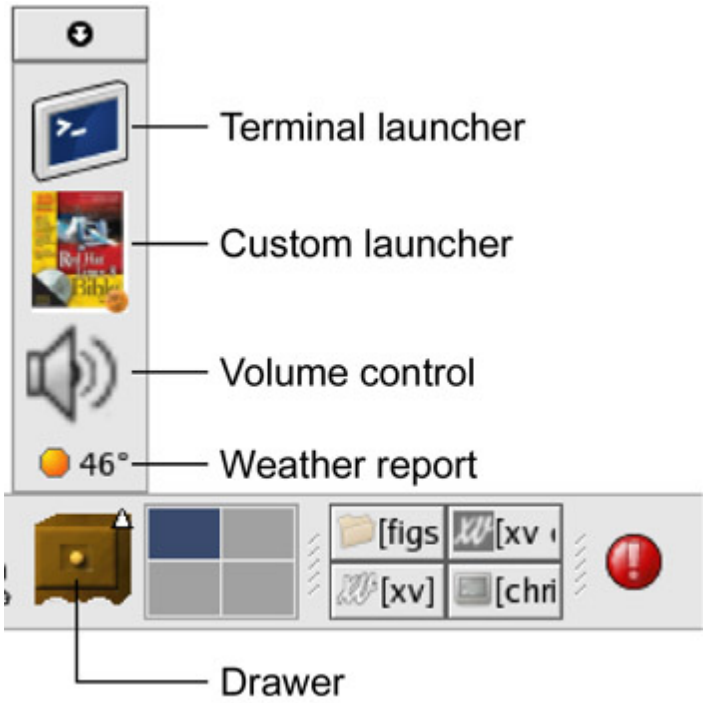

Figure 3-6: A drawer is a great way to contain personal utilities and launchers.

You can do much more with the desktop and the panel. To learn more about configuring your desktop, you can check out the specific descriptions of the GNOME and KDE desktops later in this chapter.

Tips for configuring your desktop

Now that you have experimented with a few things on the desktop, you should configure certain features to get Red Hat Linux really working well for you. Most of the tips I describe here will help you get Red Hat Linux working well on the network.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Some of the tips described here should be carried out by the system administrator. They apply to you if you are the system administrator for your organization, or if you are co nfig

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Register with Red Hat Network — If you see a blinking red exclamation point icon, it means that updates are available to your Red Hat Linux system. I strongly recommend that you register with Red Hat Network to get notifications of updates and perhaps have those updates automatically installed. These updates can patch dangerous security holes in a timely fashion. At the time of this printing, Red Hat was offering a $60/year subscription plan that I would recommend for anyone who has a Red Hat Linux computer connected all the time to a public network (such as the Internet). It is a small price to pay for piece of mind.

Cross-Reference Refer to Chapters 2 and 10 for more information on Red Hat Network and the up2date facility.

Set up your network — You may have already configured your network interfaces (dial-up or LAN) during installation. If not, refer to Chapter 15 for setting up a LAN and Chapter 16 for setting up an Internet connection.

Configure e-mail — You must identify information about your e-mail account in order to use e-mail. Click the Evolution E-mail icon in the panel to start the process of configuring e-mail. Refer to Chapter 9 for information on setting up and using e-mail.

Configure the Web browser — Open the Mozilla Web Browser from the panel. Although it should work fine at browsing the Internet once you have a network connection set up, there are a few things you should do to tune your browser. These include choosing a home page (click Edit → Preferences, then type a home page location); setting the browser window's appearance (from the Preferences window, click Appearance and select a topic); and importing bookmarks (click Bookmarks → Manage Bookmarks, then select Tools → Import from the Bookmark Manager window).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you are coming from a Windows environment, you may find that some Web content doesn't work well in Mozilla. Refer to Chapter 9 for suggestions on ways to enhance e

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The sections that follow provide more details on using the GNOME and KDE desktops.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the GNOME Desktop

GNOME (pronounced *guh-nome*) provides the desktop environment that you get by default when you install Red Hat Linux. This desktop environment provides the software that is between your X Window System framework and the look-and-feel provided by the window manager. GNOME is a stable and reliable desktop environment, with a few cool features in it.

The new GNOME 2.4 desktop comes with the most recent version of Red Hat Linux. For GNOME 2.4, enhancements include a CD-burning feature in the Nautilus file manager, improved panels, and accessibility features to improve ease-of-use for people with disabilities

To use your GNOME desktop, you should become familiar with the following components:

Metacity (Window Manager) — The default window manager for GNOME in Red Hat Linux is Metacity. The window manager provides such things as themes, window borders, and window controls.

Nautilus (File Manager/Graphical Shell) — When you open a folder (for example, by clicking the Home icon on your desktop), the Nautilus window opens and displays the contents of the selected folder. Nautilus can also display other types of content, such as shared folders from Windows computers on the network (using SMB).

GNOME Panel (Application/Task launcher) — This panel, which lines the bottom of your screen, is designed to make it convenient for you to launch the applications you use, manage running applications, and work with multiple virtual desktops. By default, the panel contains the main menu (represented by a red hat), desktop application launchers (Mozilla browser, Evolution e-mail, and a set of OpenOffice applications), a workspace switcher (for managing four virtual desktops), window list, and a clock.

Desktop Area — The windows and icons you use are arranged on the desktop area. The desktop area supports such things a drag-and-drop between applications, a desktop menu (right click to see it), and icons for launching applications.

Besides the components just described, GNOME includes a set of Preferences windows that let you configure different aspects of your desktop. You can change backgrounds, colors, fonts, keyboard shortcuts, and other features relating to the look and behavior of the desktop. Figure 3-7 shows how the GNOME desktop environment appears the first time you log in, with a few windows added to the screen.

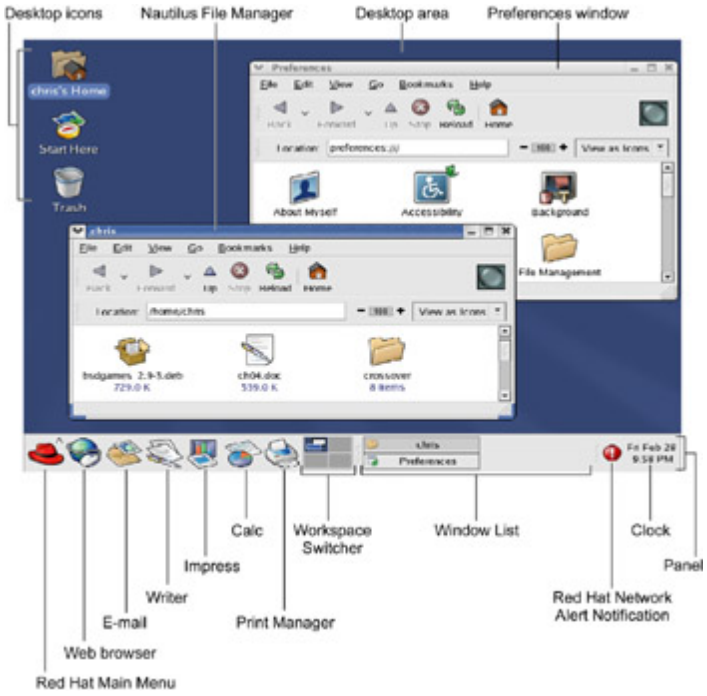

Figure 3-7: In the GNOME desktop environment, you can manage applications from the panel.

The following sections provide details on using the GNOME desktop.

Using the Metacity window manager

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The Metacity window manager seems to have been chosen as the default window manager for GNOME in Red Hat Linux because of its simplicity. The creator of Metacity refers to it as a "boring window manager for the adult in you" — then goes on to compare other window managers to colorful, sugary cereal while Metacity is characterized as Cheerios.

There really isn't much you can do with Metacity (except get your work done efficiently). Assigning new themes to Metacity, changing colors and window decorations, is done through the GNOME preferences (and is described later). A few Metacity themes exist, but expect the number to grow.
Basic Metacity functions that might interest you are keyboard shortcuts and the workspace switcher. Table 3-1 shows keyboard shortcuts to get around the Metacity window manager.

Table 3-1: Metacity Keyboard Shortcuts

Actions	Keystrokes
Window focus cycle forward, with pop-up icons	Alt+Tab
cycle backwards, with pop-up icons	Alt+Shift+Tab
cycle forward, without pop-up icons	Alt+Esc
cycle backwards, without pop-up icons	Alt+Shift+Esc
Panel focus cycle forward among panels	Alt+Ctrl+Tab
cycle backwards among panels	Alt+Ctrl+Shift+Tab
Workspace focus cycle to next workspace	Ctrl+Alt+right arrow
cycle to previous workspace	Ctrl+Alt+left arrow
Minimize/unminimize all windows	Ctrl+Alt+D
Show window menu	Alt-Spacebar
Close menu	Esc

Another Metacity feature that may interest you is the workspace switcher. Four virtual workspaces appear in the workspace switcher on the GNOME panel. Here are some things to do with the workspace switcher:

Choose current workspace — Four virtual workspaces appear in the workspace switcher. Click on any of the four virtual workspaces to make it your current workspace.
Move windows to other workspaces — Click any window, each represented by a tiny rectangle in a workspace, to drag and drop it to another workspace.
Add more workspaces — Right-click the workspace switcher, and select Preferences. You can add workspaces (up to 32).
Name workspaces — Right-click the workspace switcher and select Preferences. Click in the Workspaces pane to change names of workspaces to any names you choose.

You can view and change information about Metacity controls and settings using the gconf-editor window (type **gconf-editor** from a Terminal window). As the window says, it is not the recommended way of changing preferences. So, when possible, you should change the desktop through GNOME preferences. However, gconf-editor is a good way to see descriptions of each Metacity feature.

From the gconf-editor window, select apps → metacity. Then choose from general, global_keybindings, keybindings_commands, window_keybindings and workspace_names. Click each key to see its value, along with short and long descriptions of the key.

Using the GNOME panel

The GNOME panel is intended to be the place from which you manage your desktop. From this panel

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

you can start applications (from buttons or menus), see what programs are active, and monitor how your system is running. There are also many ways to change the panel — by adding applications or monitors, or by changing the placement or behavior of the panel, for example.
Click any open space on the panel to see the Panel menu. The Panel menu appears as shown in Figure 3-8.

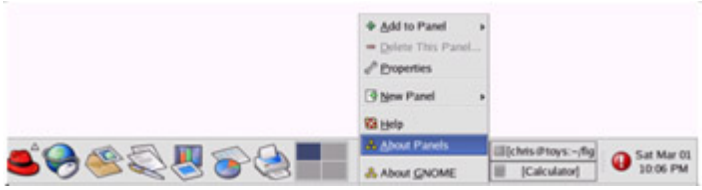

Figure 3-8: Left-click any open spot on the GNOME Panel to see the Panel menu.

From the GNOME Panel menu, you can perform a variety of functions, including:

Use the Red Hat menu — Displayed on the Red Hat menu are most of the applications and system tools you will use from the desktop.
Add to panel — You can add an applet, menu, launcher, drawer, or button.
Delete This panel — You can delete the current panel.
Properties — Change position, size, and background of the panel.
New panel — You can add panels to your desktop in different styles and locations.

There also ways of working with items on a panel. These include:

Moving — To move items on a panel, simply drag-and-drop them to a new postion.
Resizing — Some elements, such as the Window List, can be resized by clicking an edge and dragging it to the new size.
Window List — Tasks running on the desktop appear in the Window List area. Click a task to have that task minimized or maximized.

The following sections describe some things you can do with the GNOME panel.

Use the Red Hat menu

Click the Red Hat on the panel and you see categories of applications and system tools that you can select. Click the application you want to launch. To add an item to launch from the panel — and to view its properties — right-click it. There is currently no way to add or remove applications to or from this menu from the GUI in GNOME. However, you can manually add items to your GNOME menus.
To add to the red hat menu, create a .desktop file in the `/usr/share/applications` directory. The easiest way to do that is copy an existing .desktop file that is on the menu you want and modify it. For example, to add a video player to the Sound & Video menu, you could do the following (as root user):

cd /usr/share/applications

cp gnome-cd.desktop vidplay.desktop

Next use any text editor to change the contents of the vidplay.desktop file you created by adding a comment, file to execute, icon to display, and application name. After you save the changes, the new item will immediately appear on the menu (no need to restart anything).

Cross-Reference	There is a graphical way of modifying your main menu if you are using the KDE desktop. See the Menus bullet in the "Changing panel attributes" section later in this chapter for details.

Adding an applet

1. There are several small GNOME applications, called *applets*, that you can run directly on the GNOME panel. These applets can show information you may want to see on an ongoing basis or may just provide some amusement. To see what applets are available and to add applets that you want to your panel, perform the following steps:Right-click on an open space in the panel so that the panel menu appears.
2. Select Add to Panel.
3. Select an applet from one of the following categories:

Accessories — Has clock, dictionary lookup, stock ticker, and

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

weather report applets.
Amusements — Includes a fortune-telling fish and eyes that follow your mouse (geyes).
Internet — Contains an e-mail Inbox monitor and modem lights monitor.
Multimedia — Includes an applet for playing CDs and a volume control.
Utility — Includes applets for monitoring battery levels, inserting special characters, running a command line, mounting floppy drives, monitoring systems load, and performing other useful tasks.

4. Select the applet to add from the category you selected. The applet appears on the panel, ready for you to use. Figure 3-9 shows, from left to right, geyes, system monitor, CD player, stock ticker, e-mail Inbox monitor, and dictionary lookup applets.

Figure 3-9: Applets let you monitor activities, play CDs, watch your mail, or look up dictionary words.

After an applet is installed, right-click it to see what options are available. For example, select Preferences for the stock ticker, and you can add or delete stocks whose prices you want to monitor. If you don't like the applet's location, right-click it, click Move, slide the mouse until the applet is where you want it (even to another panel), and click to set its location.

If you no longer want an applet to appear on the panel, right-click it, and then click Remove From Panel. The icon representing the applet will disappear. If you find that you have run out of room on your panel, you can add a new panel to another part of the screen, as described in the next section.

Adding another panel

You can have several panels on your GNOME desktop. You can either add edge panels (that run along the entire bottom, top, or side of the screen) or a corner panel (that only expands as needed to show the applets it contains). To add a panel, do the following:

1. Right-click an open space in the panel so that the Panel menu appears.
2. Select New Panel.
3. Select from the following menu items:
   Corner panel — Adds a panel in the corner that stretches out only as far as is required to display the applets it contains.
   Edge panel — Creates a panel along the edge of the screen (top, sides, or bottom).
   Floating panel — Adds a panel that begins from the middle of the edge of a display (instead of from a corner).
   Sliding panel — Creates a panel that stretches out only as far as required to display the applets it contains. When you hide/unhide the panel, it closes in the middle of the display instead of just at the edges.
   Menu panel — Places a menu bar across the top of the screen.

After you've added a panel, you can add applets or application launchers to it as you did to the default panel. To remove a panel, right-click it and select Delete This Panel.

Adding an application launcher

Icons on your panel represent a Web browser and several office productivity applications. You can add your own icons to launch applications from the panel as well. To add a new application launcher to the panel, do the following:

1. Right-click in an open space on the panel.
2. Select Add to Panel → Launcher from the menu. All application categories from your main desktop menu (the one under the red hat) appear.
3. Select the category of application you want; then select the application. An icon

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

representing the application appears.

To launch the application you just added, single-click it.

If the application you want to launch is not on your red hat menu, you can build one yourself as follows:

1. Right-click in an open space on the panel.
2. Select Add to Panel → Launcher. The Create Launcher window appears.
3. Provide the following information for the application that you want to add:
   Name — A name to identify the application (this appears in the tool tip when you mouse is over the icon).
   Generic Name — A name to identify the type of application.
   Comment — A comment describing the application. Like with Name, this information appears when you later move your mouse over the launcher.
   Command — The command line that is run when the application is launched. You should use the full path name, plus any required options.
   Type — Select Application (to launch an application). (Other selections include Link, to open a Web address in a browser or FSDevice to open a file system.)
   Run in Terminal — If it is a character-based or ncurses application. (Applications written using the curses library run in a Terminal window but offer screen-oriented mouse and keyboard controls.)
4. Click the Icon box (it may say No Icon). Select one of the icons shown and click OK. Alternatively, you can browse the Red Hat Linux file system to choose an icon.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Icons available to represent your application are contained in the /usr/share/pixmaps directory. These icons are either in png or xpm formats. If there isn't an icon in the directory you u

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

5. Click OK.

The application should now appear in the panel. Click it to start the application.

Adding a drawer

By adding a drawer to your GNOME panel, you can add several applets and launchers and have them take up only one slot on your panel. You can use the drawer to show the applets and launchers as though they were being pulled out of a drawer icon on the panel.

To add a drawer to your panel, right-click the panel and then select Add to Panel → Drawer. The drawer should appear on the panel. The drawer behaves just like a panel. Right-click the drawer area, and add applets or launchers to it as you would to a panel. Click the drawer icon to retract the drawer. Figure 3-10 shows a portion of the panel that includes an open drawer. This example includes icons for launching a Terminal window, the GIMP, and the Ethereal window.

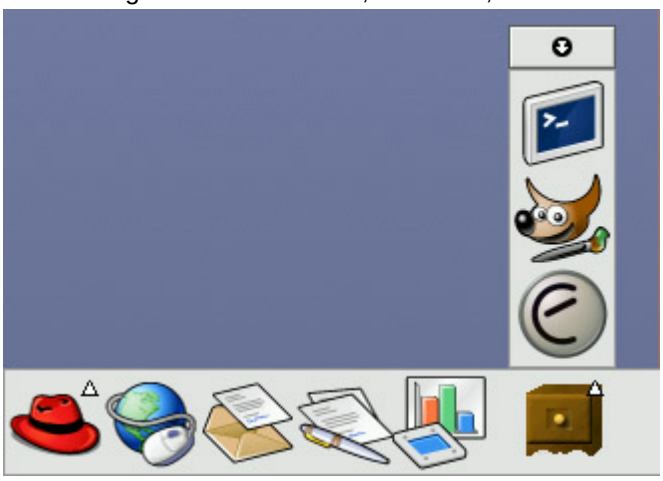

Figure 3-10: Add launchers or applets to a drawer on your GNOME panel.

Changing panel properties

Properties you can change that relate to a panel are limited to the position, size, hiding policy, and background. To open the Panel properties window that applies to a specific panel, right-click on an open space on the panel, then choose Properties. The Panel properties window that appears includes following values:

Panel Type — You can see whether the panel is an Edge, Aligned, Sliding, or Floating panel (you can't change this value).
Position — Depending on the type of panel, you can move it to different locations on the screen by clicking on a new position.
Size — You can select the size of your panel, from XX Small (12 pixels) to XX Large (128 pixels).
Autohide — You can select whether or not a panel is automatically hidden (appearing only when the mouse pointer is in the area).
Show hide buttons — You can choose whether or not the Hide/Unhide buttons (with pixmap arrows on them) appear on the edges of the panel.
Arrows on hide buttons — If you select "Show hide buttons" you can select to either have arrows on those buttons or not.
Background — From the Background tab, you can assign a color to the background of the panel, assign a pixmap image, or just leave the Default (sort of gray). Click the Background Type box, if you want to select an Image for the background, then select an image, such as a tile from `/usr/share/backgrounds/tiles` or other directory.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip I usuall y turn on the AutoHi de featur e and turn off the Hide button s. Using AutoHi de gives you more space to work with on your deskt op. When you move your mouse to the edge

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the Nautilus file manager

At one time, file managers did little more than let you run applications, create data files, and open folders. These days, as the information a user needs expands beyond the local system, file managers are expected to also display Web pages, access FTP sites, and play multimedia content. The Nautilus file manager, which is the default GNOME file manager, is an example of just such a file manager. When you open the Nautilus file manager window (from the GNOME main menu or by opening the Home icon or other folder on your desktop), you see the name of the location you are viewing (such as the folder name) and what that location contains (files, folders, and applications). Figure 3-11 is an example of the file manager window displaying the home directory of a user named chris (`/home/chris`).

Figure 3-11: Move around the file system, open directories, launch applications, and open Samba folders.

Icons on the toolbar of the Nautilus window let you move forward and back among the directories and Web sites you visit. To move up the directory structure, click the up arrow. To refresh the view of the folder or Web page, click the Refresh button. The Home button takes you to your home page and the Web search button lets you search the Web.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note One note about viewing HTML content (in other words, Web pages). The version of Nautilus that originally came with GNOME 2.4 did not properly display HTML content. Depending on how

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Icons in Nautilus often indicate the type of data that a particular file contains. The contents or file extension of each file can determine which application is used to work with the file. Or, you can right-click an icon to open the file it represents with a particular application or viewer.

Some of the more interesting features of Nautilus are described below:

Sidebar — Click on View → Side Pane to have a sidebar appear in the left column of the screen. From the sidebar, you can click on tabs that represent different types of information you can select. The Tree tab shows a tree view of the directory structure, so you can easily traverse your directories.

The Notes tab lets you add notes that become associated with the current Directory or Web page. The History tab displays a history of directories and Web sites you have visited, allowing you to click those items to return to the sites they represent. Right-click in the sidebar to choose which of the sidebar tabs are displayed.

Windows File and Printer Sharing — If your computer is connected to a LAN on which Windows computers are sharing files and printers, you can view those resources from Nautilus. Type **smb:** in the location box to see available workgroups. Click a workgroup to see computers from that workgroup that are sharing files and printers. Figure 3-12 shows an example of Nautilus displaying icons representing Windows computers in a workgroup called estreet (`smb://estreet`).

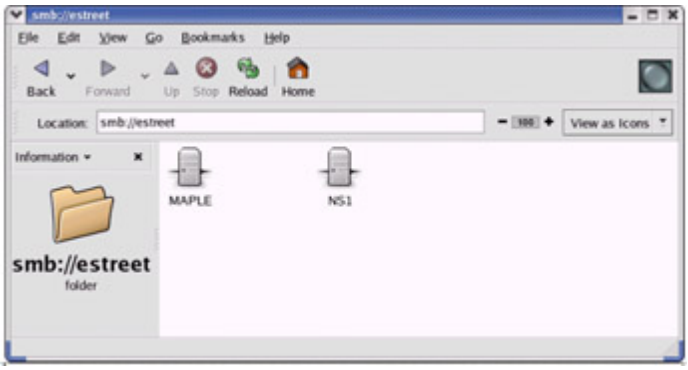

Figure 3-12: Display shared Windows file and printer servers (SMB) in Nautilus.

MIME types and file types — To handle different types of content that may be encountered in the Nautilus window, you can set applications to respond based on MIME type and file type. With a folder being displayed, right-click a file for which you want to assign an application. Click either Open With an Application or Open With a Viewer. If no application or viewer has been assigned for the file type, click Associate Application to be able to select an application. From the Add File Types window, you can add an application based on the file extension and MIME type representing the file.

Cross-Reference For more information in MIME types, see the description of MIME types in the "Changing GNOME preferences" section later in this chapter.

Drag-and-Drop — You can use drag-and-drop within the Nautilus window, between the Nautilus and the desktop, or between multiple Nautilus windows. As other GNOME-compliant applications become available, they are expected to also support the GNOME drag-and-drop feature.

If you need more information on the Nautilus file manager, visit the GNOME Web site (`www.gnome.org/nautilus`).

Changing GNOME preferences

There are many ways to change the behavior, look, and feel of your GNOME desktop. Most GNOME preferences can be modified from windows you can launch from the main red hat menu or from the Preferences folder. You can open that folder from the red hat menu by clicking Preferences → Control Center (or typing **preferences:** in the Nautilus location box).

Unlike earlier versions of GNOME for Red Hat Linux, boundaries between preferences relating to the window manager (Metacity), file manager (Nautilus), and the GNOME desktop itself have been blurred. Preferences for all of these features are in the Preferences window. Figure 3-13 shows the Preferences window, with icons that represent features you can change.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

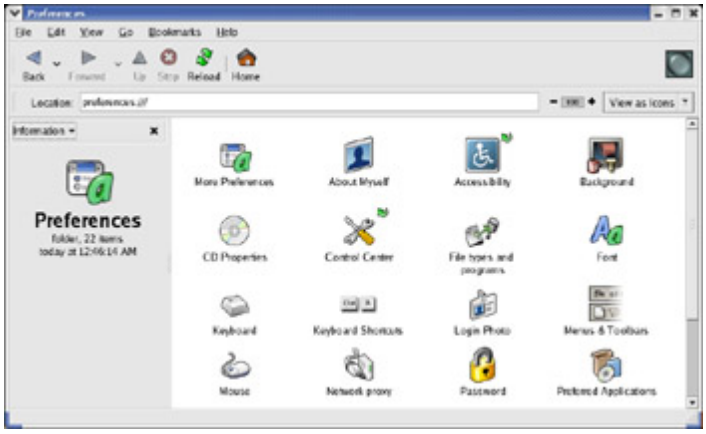

Figure 3-13: Change the look-and-feel of your desktop from the Preferences window.

The following items highlight some of the preferences you might want to change:

Accessibility — If you have difficulty operating a mouse or keyboard, the Keyboard Accessibility Configuration window lets you adapt mouse and keyboard settings to make those devices more accessible. From the Preferences window, open Accessibility.

Background — From Background Preferences, you can choose a single solid color, a two-color gradient (shading from top to bottom or left to right), or an image to use as wallpaper. If you choose a solid color (under Background Style), click the Color box, choose a color from the palette, and select OK. To have the color fade into another color, click Vertical or Horizontal Gradient, and then select the second color. Choosing Vertical has the colors transform from top to bottom, while Horizontal causes the colors to change from left to right.

To use wallpaper for your background, open a folder containing the image you want to use. Then drag the image into the Background Preferences window. You can choose from a variety of images in the `/usr/share/nautilus/patterns` and `/usr/share/backgrounds/tiles` directories. Then, choose to have the image as wallpaper (repeated pattern), centered, scaled (in proportion), or stretched (using any proportion to fill the screen).

CD Properties — Even if you don't change CD properties, it is important to know what happens when you insert a Data CD, Audio CD, Blank CDs, or DVD (video) disk. These properties are associated with a feature called "magicdev, " which is a bit controvertial.

For Data CDs , the CD is mounted when it is inserted, any auto-run program on the CD is run, and a file manager window opens for the CD you inserted. If you would rather mount and open the CD as you choose, you can turn off any or all of these preferences.

For Audio CDs, the gnome-cd player is launched and the CD begins playing. You can type in a different CD player, if you like, or turn off "Run command when CD is inserted" so that you can choose which player to use later.

For Blank CDs, a CD-burning utility is launched. After that, you can burn audio files or data to the blank CD.

For DVD (Video), the DVD is not set to play automatically. If you have a player installed that can play the content of DVDs that you have, turn on this feature and add the command to run the player into the Command box.

File Types and Programs — The File Types and Programs preferences can help you understand the different types of data files that GNOME knows about. Double-click this icon to see data types (audio, documents, images, information, and so on) that have definitions in GNOME. Then choose a particular data type (such as Audio, ogg audio).

From the Edit file type window that appears, you can see the information assigned to the file type. So, for example, when data that ends with a .ogg extension appears in a Nautilus

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

window, you can see the icon that will represent the file, the mime type assigned to the file, and the action (if any) is taken when you open the file.

You can modify any file type that appears in these preferences windows. You can choose what applications are run and and what icons represent data of this type. You can even create your own data types.

Screensaver — You can choose from dozens of screensavers from the Screensaver Preferences window. Select Random Screensaver to have your screen saver chosen randomly, or select one that you like from the list. Next, choose how long your screen must be idle before the screensaver starts (default is 10 minutes). You can also choose to require a password or to enable power management to shut down your monitor after a set number of minutes (Advanced Tab). Figure 3-14 shows Screensaver preferences.

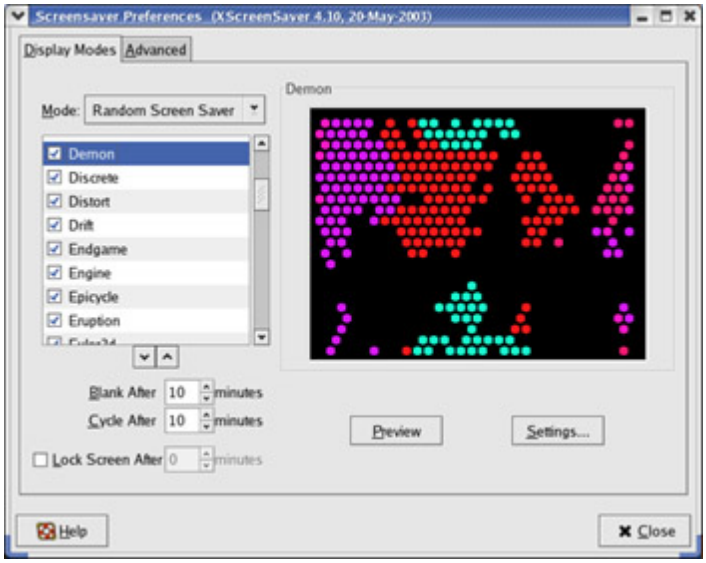

Figure 3-14: Select specific or random screen savers from Screensaver Preferences.

Theme Selector — You can choose to have an entire theme of elements be used on your desktop. A desktop theme affects not only the background, but also the way that many buttons and menu selections appear. There are only a few themes available for the window manager (Metacity) in the Red Hat Linux distribution. You can get a bunch of other Metacity themes from `themes.freshmeat.net` (click on Metacity).

Click Install theme; then click the Window Border tab to select from different themes that change the title bar and other borders of your windows. Click the Icons tab to choose different icons to represent items on your desktop. Themes change immediately as you click or when you drag a theme name on the desktop.

Exiting GNOME

When you are done with your work, you can either log out from your current session or shut down your computer completely. To exit from GNOME, do the following:

1. Click the Red Hat menu button.
2. Select Log Out from the menu. A pop-up window should appear, asking if you want to log out, shut down, or restart the computer.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip At this point, you can also select to save your session by clicking Save current setup. This is a great way to have the applications that you use all the time restart the next tim

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

3. Select Logout from the pop-up menu. This will log you out and return you to either the graphical login screen or to your shell login prompt. If, instead, you select Shut Down, the system is shut down; if you select Reboot, the system is restarted.
4. Select OK to finish exiting from GNOME.

If you are unable to get to the Log out button (if, for example, your Pager was deleted by mistake), there are two other exit methods. Try one of these two ways, depending on how you started the desktop:

If you started the desktop by typing **startx** from your login shell, press Ctrl+Alt+F1 to return to your login shell. Then type Ctrl+C to kill the desktop.

If you started the desktop from a graphical login screen, first open a Terminal window (right-click the desktop and then select New Terminal). In the Terminal window, type **ps x | more** to see a list of running processes. Look for a command named `GNOME-session` and determine its number under the PID column. Then type **kill *PID*** , where *PID* is replaced by the PID number. You should see the graphical login screen.

Although these are not the most graceful ways to exit the desktop, they work. You should be able to log in again and restart the desktop.

Team LiB

PREVIOUS NEXT

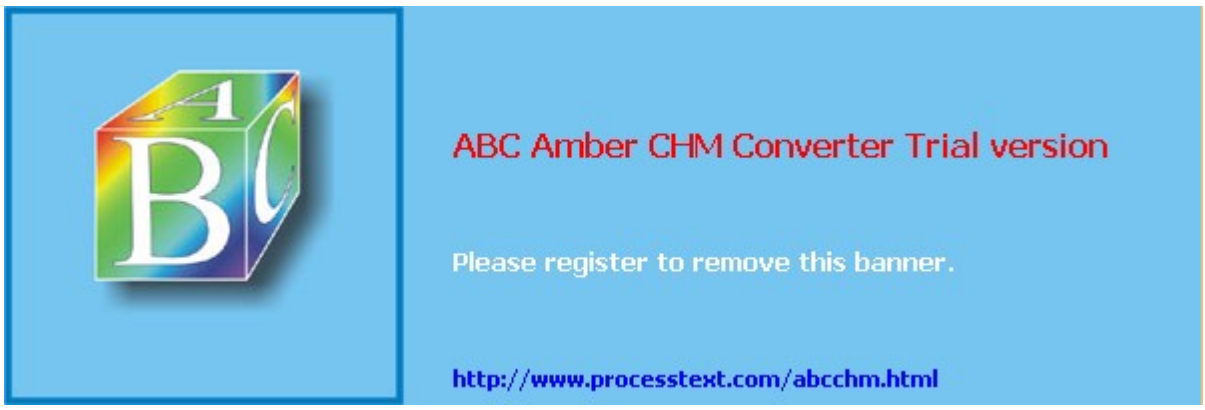

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the KDE Desktop

The KDE desktop was developed to provide an interface to Linux and other UNIX systems that could compete with MacOS or Microsoft Windows operating systems for ease of use. Integrated within KDE are tools for managing files, windows, multiple desktops, and applications. If you can work a mouse, you can learn to navigate the KDE desktop. Red Hat Linux includes version 3.1.2 of KDE.

Note KDE is not installed by default for Red Hat Linux personal desktop and workstation installations. Therefore, to use the procedures in this section, you might have to install KDE. See Chapter 2 for information on how to use the Add and Remove Software window so you can add KDE.

The lack of an integrated, standardized desktop environment in the past has held back Linux and other UNIX systems from acceptance on the desktop. While individual applications could run well, you rarely could drag-and-drop files or other items between applications. Likewise, you couldn't open a file and expect the machine to launch the correct application to deal with it or save your windows from one login session to the next. KDE provides a platform for developers to create programs that easily share information and detect how to deal with different data types.

The following section describes how to get started with KDE. This includes using the KDE Setup wizard, maneuvering around the desktop, managing files, windows, virtual desktops, and adding application launchers.

Starting with KDE

You can select the KDE desktop from the login screen (provided that KDE is installed). Choose Session → KDE. Then type your login name and password, as prompted. The KDE desktop should appear, as shown in Figure 3-15.

Figure 3-15: Manage files and applications graphically with the KDE desktop.

KDE desktop described

You may notice that, in Red Hat Linux, KDE looks very similar to GNOME. That's because both begin with the same icons and background. Beneath the surface, however, you will notice that there are a lot more tools with KDE for configuring and using your desktop.

Here are some descriptions of what you will find on the KDE desktop for Red Hat Linux:

Desktop icons — The desktop starts with several icons on it to provide quick access

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

to selected features. By default, you will probably have an icon that gives you access to your home directory. You will also see icons representing your removable drives (CDs, DVDs, and floppies), the Start Here icon (to open applications), and the Trash.
Panel — The panel provides some quick tools for launching applications and managing the desktop. You can adapt the panel to your needs by resizing it, adding tools, and changing the look-and-feel. By default, you start with the same icons you had in GNOME: one for your Web browser and several for office-productivity applications.**KDE Menu** — This panel button is represented by a K icon. Click this button to see a menu of applications, utilities, and configuration tools that are available to run on your KDE desktop.
Web Browser — This panel button looks like a globe with a mouse on it. Click it to open the Mozilla Web Browser.
Evolution Email — This panel button is the image of an envelope with a letter on it. Click it to open the Evolution window for sending, receiving, and managing e-mail.
Writer — This button, which includes a pen and paper, launches the OpenOffice Writer word-processing window.
Impress —This button includes a bar chart image and launches the OpenOffice Impress presentation window. Impress lets you create presentations similar to those produced by Microsoft PowerPoint.
Calc — This button includes an image of a pie chart and launches the OpenOffice Calc spreadsheet application.
Print Manager — This button opens a Print Manager window to manage printers and view queued print jobs.
Desktop Pager — This box on the panel consists of your virtual desktops, which contain small views of each desktop. There are four virtual desktops available to you, by default. These are labeled 1, 2, 3, and 4. You begin your KDE session on virtual desktop 1. You can change to any of the four desktops by clicking it.
Taskbar — This button shows the tasks that are currently running on the desktop. The window that is currently active appears pressed in. Click a task to toggle between opening and minimizing the window.
Klipper — This button looks like a clipboard. Click it if you want to clear your desktop's clipboard.
Red Hat Network Alert Notification Tool — When software updates are available to Red Hat Linux, this changes from a check mark to a blinking, red exclamation mark. Click it to register with Red Hat Network and to get software updates.
Clock — The current time and date are shown on the far right-hand side of the panel. Click them to see a calendar for the current month. Click the arrow keys on the calendar to move forward and back to other months.
Konqueror — The default file manager for KDE.
KDE Control Center — The window for changing desktop preferences.

Getting around the desktop

Navigating the desktop is done with your mouse and keyboard. You can use a two-button or three-button mouse. Using the keyboard to navigate requires some Alt and Ctrl key sequences.

Using the mouse

The responses from the desktop to your mouse depend on which button you press and where the mouse pointer is located. Table 3-2 shows the results of clicking each mouse button with the mouse pointer placed in different locations.

Table 3-2: Mouse Actions

Pointer Position	Mouse Button	Results
Window title bar or frame (current window active)	Left	Raise current window.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 3-2: Mouse Actions

Pointer Position	Mouse Button	Results
Window title bar or frame (current window active)	Middle	Lower current window.
Window title bar or frame (current window active)	Right	Open operations menu.
Window title bar or frame (current window not active)	Left	Activate current window and raise it to the top.
Window title bar or frame (current window not active)	Middle	Activate current window and lower it.
Window title bar or frame (current window not active)	Right	Activate current window without changing position.
Inner window (current window not active)	Left	Activate current window, raise it to the top, and pass the click to the window.
Inner window (current window not active)	Middle	Activate current window and pass the click to the window.
Inner window (current window not active)	Right	Activate current window and pass the click to the window.
Any part of a window	Left (plus hold Alt key)	Move the window to a new location. This is very useful if the title bar of a window is off the screen. With this action, you can grab any part of the window to move it.
Any part of a window	Middle (plus hold Alt key)	Toggle between raising and lowering the window.
Any part of a window	Right (plus hold Alt key)	Resize the window.
On the desktop area	Left (hold and drag)	Select a group of icons.
On the desktop area	Right	Open system pop-up menu.

The mouse actions in the table are all single-click actions. Use single-click with the left mouse button to open an icon on the desktop. On a window title bar, double-clicking results in a window-shade action, where the window scrolls up and down into the title bar.

Using keystrokes

If you don't happen to have a mouse or you just like to keep your hands on the keyboard, there are

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

several keystroke sequences you can use to navigate the desktop. Here are some examples:

Step through desktops (Ctrl+Tab) — To go from one virtual desktop to the next, hold down the Ctrl key and press the Tab key until you see the desktop that you want to make current. Then release the Ctrl key to select that desktop.

Step through windows (Alt+Tab) — To step through each of the windows that are running on the current desktop, hold down the Alt key and press the Tab key until you see the one you want. Then release the Alt key to select it.

Open Run Command box (Alt+F2) — To open a box on the desktop that lets you type in a command and run it, hold the Alt key and press F2. Next, type the command in the box and press Enter to run it. You can also type a URL into this box to view a Web page.

Close the current window (Alt+F4) — To close the current window, press Alt+F4.

Close another window (Ctrl+Alt+Esc) — To close an open window on the desktop, press Ctrl+Alt+Esc. When a skull and cross bones appears as the pointer, move the pointer over the window you want to close and click the left mouse button. (This is a good technique for killing a window that has no borders or menu.)

Switch virtual desktops (Ctrl+F1, F2, F3 or F4 key) — To step through virtual desktops, press and hold the Ctrl key and press F1, F2, F3, or F4 to go directly to desktop one, two, three, or four, respectively. You could do this for up to eight desktops, if you have that many configured.

Open window operation menu (Alt+F3) — To open the operations menu for the active window, press Alt+F3. When the menu appears, move the arrow keys to select an action (Move, Size, Minimize, Maximize, etc.), then press Enter to select it.

Managing files with the Konqueror File Manager

The Konqueror File Manager helps elevate the KDE environment from just another X window manager to an integrated desktop that can compete with GUIs from Apple Computing or Microsoft. The features in Konqueror rival those that are offered by those user-friendly desktop systems. Figure 3-16 shows an example of the Konqueror File Manager window.

Figure 3-16: Konqueror provides a network-ready tool for managing files.

Some of Konqueror's greatest strengths over earlier file managers are the following:

Network desktop — If your computer is connected to the Internet or a LAN, features built into Konqueror let you create links to files (using FTP) and Web pages (using HTTP) on the network and open them within the Konqueror window. Those links can appear as file icons in a Konqueror window or on the desktop. When a link is opened (single-click), the contents of the FTP site or Web page appears right in the Konqueror window.

Web browser interface — The Konqueror interface works like Mozilla, Internet Explorer, or another Web browser in the way you select files, directories, and Web content. A single-click opens a file, link to a network resource, or application program. You can also open content by typing Web-style addresses in a Location box.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Web pages that contain Java content will not run by default in Konqueror. To turn on Java support, choose Settings → Configure Konqueror. From the Settings window, click Java & JavaS

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

File types and MIME types — If you want a particular type of file to always be launched by a particular application, you can configure that file yourself. KDE already has dozens of MIME types defined that can automatically detect particular file and data types and start the right application. There are MIME types defined for audio, image, text, video, and a variety of other content types.

Of course, you can also perform many standard file manager functions with Konqueror. For manipulating files, you can use features like Select, Move, Cut, Paste, and Delete. You can search directories for files, create new items (files, folders, and links, to name a few), view histories of the files and Web sites you have opened, and create bookmarks.

Working with files

Because most of the ways of working with files in Konqueror are quite intuitive (by intention), I'll just give a quick rundown of how to do basic file manipulation:

Open a file — Double-click the left mouse button on a file. The file will open right in the Konqueror window, if possible, or in the default application set for the file type. You can also open a directory (to make it the current directory), application (to start the application), or link (to open the target of a link) in this way.
Choose an application — Single-click the left mouse button to highlight it or right mouse button to open a menu. When you right-click a data file, select the Open With menu. The menu that appears shows which applications are set up to open the file.
Delete a file — Single-click the right mouse button and select Delete. You are asked if you really want to delete the file. Click Yes to permanently delete it.
Copy a file — Single-click the right mouse button and select Copy. This copies the file to your clipboard. After that, you can paste it to another folder. (See the "Move a file" bullet for a drag-and-drop method of copying.)
Paste a file — Single-click the right mouse button (on an open area of a folder) and select Paste. A copy of the file you copied previously is pasted in the current folder.
Move a file — With the original folder and target folder both open on the desktop, press and hold the left mouse button on the file you want to move, drag the file to an open area of the new folder, and release the mouse button. From the menu that appears, click Move. (You could also copy or create a link to the file using this menu.)
Link a file — Drag-and-drop a file from one folder to another. When the menu appears, click Link Here. (A linked file lets you access a file from a new location without having to make a copy of the original file. When you open the link, a pointer to the original file causes it to open.)

There are also several features for viewing information about the files and folders in your Konqueror windows:

View quick file information — Position the mouse pointer over the file. When a mouse pointer is over a file in a Konqueror window, information appears in the window footer. This includes the filename, file size, and file type.
View hidden files — Select View → Show Hidden Files. This allows you to see files that begin with a dot (.). Dot files tend to be used for configuration and don't generally need to be viewed in your daily work.
View file system tree — Select View → View Mode → Tree View. This presents a tree view of your folder, displaying folders above the current folder in the file system. You can click a folder in the tree view to jump directly to that folder. There are also Multicolumn, Detailed List, and Text views available.
Change icon view — Select View → Icon Size, and then choose Large, Medium, or Small to select the size of the icons that are displayed on the window. You can also choose Default Size, to return to the default icon size.

To act on a group of files at the same time, there are a couple of actions you can take. Choose Edit → Select. A pop-up window lets you match all (*) or any group of documents indicated by typing letters, numbers, and wildcard characters. Or, to select a group of files, left-click in an open area of the folder and drag the pointer across the files you want to select. All files within the box will be highlighted. When

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

files are highlighted, you can move, copy, or delete the files as described earlier.

Searching for files

If you are looking for a particular file or folder, you can use the Konqueror Find feature. To open a Find window to search for a file, choose Tools → Find File. (You could also start the kfind window by typing **kfind** from a Terminal window.) Figure 3-17 shows the kfind window.

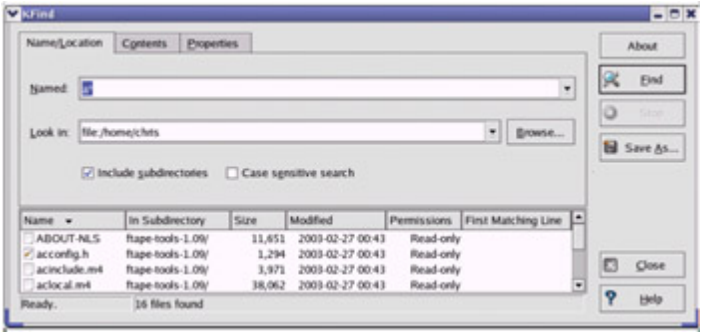

Figure 3-17: Search for files and folders from the kfind window.

Simply type the name of the file you want to search for (in the Named box) and the folder, including all subfolders, you want to search in (in the Look in box). Then click the Find button. Use metacharacters, if you like, with your search. For example, search for `*.rpm` to find all files that end in `.rpm` or `z*.doc` to find all files that begin with `z` and end with `.doc`. You can also select to have the search be case-sensitive or click the Help button to get more information on searching.

To further limit your search, you can click the Date Range tab, then enter a date range (between), a number of months before today (during the previous *x* months), or the number of days before today (during the previous *x* days). Select the Advanced tab to choose to limit the search to files of a particular type (of Type), files that include text that you enter (Containing Text), or that are of a certain size (Size is) in kilobytes.

Creating new files and folders

You can create a variety of file types when using the Konqueror window. Choose Edit → Create New, and select one of the following types:

Directory — Opens a box that lets you type in a new directory name. Type the name, click OK, and an icon representing the new directory appears.

Illustration Document — Opens a box that lets you create a document in kontour format (an illustration). Type the document name you want to create and click OK. The document should have a .kil suffix if you want it to automatically open in kontour.

HTML File — Opens a box that lets you type the name of an HTML file to create.

Text File — Opens a box that lets you create a document in text format and place it in the Konqueror window. Type the name of the text document to create and click OK.

CD/DVD-ROM Device — Opens a box that lets you type a new CD-ROM or DVD-ROM device name. Click the Device tab and type the device name (`/dev/cdrom`), the mount point (such as `/mnt/cdrom`), and the file system type (you can use iso9660 for the standard CD-ROM file system, ext2 for Linux, or msdos for DOS). When the icon appears, you can open it to mount the CD-ROM or DVD-ROM and display its contents.

Floppy Device — Opens a box to type a new floppy name. Click the Device tab and type the device name (`/dev/fd0`), the mount point (such as `/mnt/floppy`), and the file system type (you can use auto to autodetect the contents, ext2 for Linux, or msdos for DOS). When the icon appears, open it to mount the floppy and display its contents.

Hard Disc — Opens a box that lets you type the name of a new hard disk or hard-disk partition. Click the Device tab and type the device (`/dev/hda1`), the mount point (such as `/mnt/win`), and the file system type (you can use auto to autodetect the contents, ext2 or ext3 for Linux, or vfat for a Windows file system). When the icon appears, you can open it to mount the file system and display its

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

contents.

Link to Application — Opens a window that lets you type the name of an application. Click the Permissions tab to set file permissions (Exec must be on if you want to run the file as an application). Click the Execute tab and type the name of the program to run (in the field: Execute on click) and a title to appear in the title bar of the application (in the field: Window Title). If it is a text-based command, select the Run in terminal box. Click the box to Run as a different user and add the user name. Click the Application tab to assign the application to handle files of particular MIME types. Click OK.

Link to Location (URL) — Selecting this menu item opens a box that lets you create a link to a Web address. Type a name to represent the address and and type the name of the URL (Web address) for the site. (Be sure to add the `http://` , `ftp://`, or other prefix.)

Presentation Document — Opens a box to create a document in kpresenter format (a presentation). Type the document name you want to create and click OK. The document should have a .kpr or .kpt suffix if you want it to automatically open in kpresenter.

Spread Sheet Document — Opens a box that lets you create a document in kspread format (a spreadsheet). Type the document name you want to create and click OK. The document should have a .ksp suffix if you want it to automatically open in kspread.

Text Document — Opens a box that enables you to create a text document in KWord. Type a filename for the text file and click OK. The document should have a .txt, .kwd, or .kwt suffix if you want it to open automatically in Kword.

Creating MIME types and applications are described later in this chapter.

Using other browser features

Because Konqueror performs like a Web browser as well as a file manager, it includes several other browser features. For example, you can keep a bookmark list of Web sites you have visited, using the bookmarks feature. Any bookmarks that you add to your bookmarks list show up in the pull-down menu that appears when you click Bookmarks. Select from that list to return to a site. There are several ways to add and change your bookmarks list:

Add Bookmark — To add the address of the page that is currently being displayed to your bookmark list, choose Bookmarks → Add Bookmark. The next time you click Bookmarks, you will see the bookmark you just added on the Bookmarks menu. In addition to Web addresses, you can also bookmark any file or folder.

Edit Bookmarks —Selecting Bookmarks → Edit Bookmarks opens a tree view of your bookmarks. From the Bookmark Editor window that appears, you can change the URLs, the icon, or other features of the bookmark. There is also a nice feature that lets you check the status of the bookmark (that is, the address available).

New Bookmark Folder — You can add a new folder of bookmarks to your Konqueror bookmarks list. To create a bookmarks folder, choose Bookmarks → New Folder. Then type a name for the new Bookmarks folder and click OK. The new bookmark folder appears on your bookmarks menu. You can add the current location to that folder by clicking on the folder name and selecting Add Bookmark.

Configuring Konqueror options

You can change many of the visual attributes of the Konqueror window. You can select which menu bars and toolbars appear. You can have any of the following bars appear on the Konqueror window: Menu bar, Toolbar, Extra Toolbar, Location Toolbar, Bookmark Toolbar. Select Settings and then click the menu item for the bar you want to have appear (or not appear). The bar appears when the checkmark is shown next to it.

You can modify a variety of options for Konqueror by choosing Settings → Configure Konqueror. The Konqueror Settings window appears, offering the following options:

Behavior (File) — Changes file manager behavior.

Appearance — Changes file manager fonts and colors.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Previews — An icon in a Konqueror folder can be made to ressemble the contents of the file it represents. For example, if the file is a JPEG image, the icon representing the file could be a small version of that image. Using the Previews features, you can limit the size of the file used (1MB is the default) since many massive files could take too long to refresh on the screen. You can also select to have any thumbnail embedded in a file to be used as the icon or have the size of the icon reflect the shape of the image used.
File Associations — Describes which programs to launch for each file type.
Behavior (Browser) — Click the Behavior (Browser) button to open a window to configure the Web browser features of Konqueror. By enabling Form Completion, Konqueror can save form data you type and, at a later time, fill that information into other forms. If your computer has limited resources, you can speed up page display by unselecting "Automatically load images" or by disabling animations.
Java and JavaScript — Use this selection to enable or disable Java and JavaScript content contained in Web pages in your Konqueror window.
Fonts — Choose which fonts to use, by default, for various fonts needed on Web pages (standard font, fixed font, serif font, sans serif font, cursive font, and fantasy font). The serif fonts are typically used in body text, while san serif fonts are often used in headlines. You can also set the Minimum and Medium font sizes.
Web Shortcuts — Click the Web Shortcut***s*** button to see a list of keyword shortcuts you can use to go to different Internet sites. For example, follow the word "ask" with a search string to search the Ask Jeeves (`www.ask.com`) Web site. (This feature doesn't appear to be working at the moment.)
History Sidebar — Click here to modify the behavior of the list of sites you have visited (the history). By default, the most recent 500 URLs are stored, and after 90 days, a URL is dropped from the list. You will also find a button to clear your history. (To view your history list in Konqueror, open the left side panel, then click on the tiny scroll icon.)
Cookies — Click the Cookies button to select whether or not cookies are enabled in Konqueror. By default, you are asked to confirm that it is okay each time a Web site tries to create or modify a cookie. You can change that to either accept or reject all cookies. You can also set policies for acceptance or rejection of cookies based on host and domain names.
Cache — Click the Cache button to indicate how much space on your hard disk can be used to store the sites you have visited (based on the value in the Disk Cache Size field).
Proxy — Click the Proxy button if you are accessing the Internet through a proxy server. You need to enter the address and port number of the computer providing HTTP and/or FTP proxy services.

Creating an Image Gallery with Konqueror

There's a neat feature in Konqueror that lets you create a quick image gallery. The feature takes a directory of images, creates thumbnails for each one, and generates an HTML (Web) page. The HTML page includes a title you choose, all image thumbnails arranged on a page, and links to the larger images. Here's how you do it:

Add images you want in your gallery to any folder (for example, `/home/jake/images`). Make sure they are sized, rotated, and cropped the way you like before beginning. (Try The Gimp for manipulating your images by typing **gimp&** from a Terminal.)
Open the folder in Konqueror (for example, type **`/home/jake/images`** in the Location box).
Click Tools → Create Image Gallery. The Create Image Gallery window appears.
Type a title for the image gallery into the Page Title box. You can also select other attributes of the gallery, such as the number of rows, information about the image to appear on the page (name, size, and dimension), the fonts, and the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

colors to use.
Click OK.

Konqueror generates the thumbnails and adds them to the thumbs directory. The image gallery page itself opens and is saved to the `images.html` file. You can now copy the entire contents of this directory to a Web server and publish your pictures on the Internet. Figure 3-18 shows an example of a Konqueror image gallery.

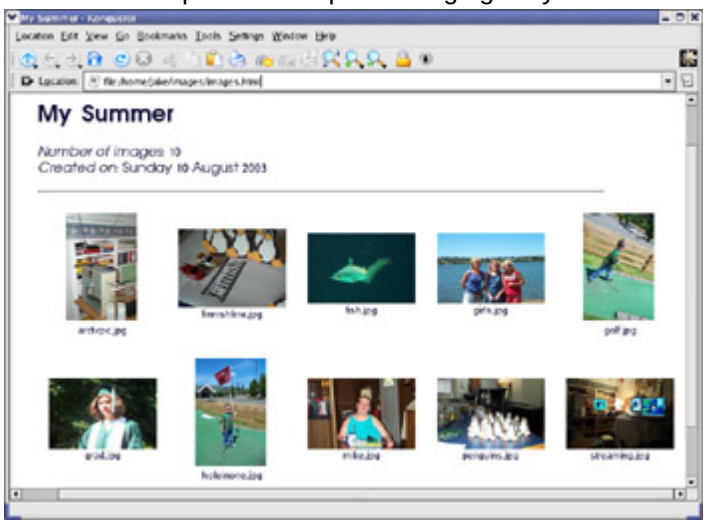

Figure 3-18: Create an image gallery in Konqueror.

Stylesheets — Click the Stylesheets button to select whether to use the default stylesheet, a user-defined stylesheet, or a custom stylesheet. The stylesheet sets the font family, font sizes, and colors that are applied to Web pages. (This won't change particular font requests made by the Web page.) If you select a custom stylesheet, click the Customize tab to customize your own fonts and colors.
Crypto — Click the Crypto button to display a list of secure certificates that can be accepted by the Konqueror browser. By default, Secure Socket Layer (SSL) version 2 and 3 certificates are accepted, as is TLS support (if supported by the server). You can also select to be notified when you are entering or leaving a secure Web site.
Browser Identification — Click the Browser Identification button to set how Konqueror identifies itself when it accesses a Web site. By default, Konqueror tells the Web site that it is the Mozilla Web browser. You can select Konqueror to appear as different Web browsers to specific sites. You must sometimes do this when a site denies you access because you do not have a specific type of browser (even though Konqueror may be fully capable of displaying the content).
Plugins — Click the Plugins button to see a list of directories that Konqueror will look in to find plug-ins. Konqueror can also scan your computer to find plug-ins that are installed for other browsers in other locations.

Managing windows

If you have a lot of windows open at the same time, tricks for organizing and managing the windows on your desktop are very helpful. KDE helps you out by maintaining window lists you can work with and shortcuts for keeping the windows in order.

Using the taskbar

When you open a window, a button representing the window appears in the taskbar at the bottom of the screen. Here is how you can manage windows from the taskbar:

Toggle windows — You can left-click any running task in the taskbar to toggle between opening the window and minimizing it.
Move windows — You can move a window from the current desktop to any other virtual desktop. Right-click any task in the taskbar, select To Desktop, then select any desktop number. The window moves to that desktop.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

All the windows that are running, regardless of which virtual desktop you are on, appear in the taskbar. If there are multiple windows of the same type shown as a single task, you can right-click that task; then select "All to Desktop" to move all related windows to the desktop you pick.

Uncluttering the desktop

If your windows are scattered willy-nilly all over the desktop, here are a couple of ways you can make your desktop's appearance a little neater:

Unclutter windows — Right-click the desktop, and then click Windows → Unclutter Windows on the menu. All windows that are currently displayed on the desktop are lined up along the left side of the screen (or aligned with other windows), from the top down.

Cascade windows — Right-click the desktop, and then click Windows → Cascade windows on the menu. The windows are aligned as they are with the Uncluttered selection, except that the windows are each indented starting from the upper-left corner.

Moving windows

The easiest way to move a window from one location to another is to place the pointer on the window's title bar; while holding down the mouse button, move the mouse so the window goes to a new location, and release the mouse button to drop the window. Another way to do it is to click the window menu button (top left corner of the title bar), click Move, move the mouse to relocate the window, and then click again to place it.

If somehow the window gets stuck in a location where the title bar is off the screen, there is a way you can move it back to where you want it. Hold down the Alt key and press the left mouse button in the inner window. Then move the window where you want it and release.

Resizing windows

To resize a window, place the pointer over a corner or side of the window border, and, while holding down the mouse button, move it until it is the size you want. Grabbing a corner lets you resize vertically and horizontally at the same time. Grabbing the side lets you resize in only one direction.

You can also resize a window from the window menu button. Click the window menu button (top left corner of the title bar) and select Size. Move the mouse until the window is resized and click to leave it there.

Using virtual desktops

To give you more space to run applications than will fit on your physical screen, KDE gives you access to several virtual desktops at the same time. Using the 1, 2, 3, and 4 buttons on the Panel, you can easily move between the different desktops. Just click the one you want.

If you want to move an application from one desktop to another, you can do so from the window menu. Click the window menu button for the window you want to move, click To Desktop, then select Desktop 1, 2, 3, or 4. The window will disappear from the current desktop and move to the one you selected.

Configuring the desktop

If you want to change the look, feel, or behavior of your KDE desktop, the best place to start is the KDE Control Center. The KDE Control Center window lets you configure dozens of attributes associated with colors, fonts, backgrounds, and screen savers used by KDE. You can also change attributes relating to how you work with windows and files.

To open the KDE Control Center from the desktop, click the K menu button and select Control Center. The KDE Control Center window appears, as shown in Figure 3-19.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Figure 3-19: Configure your desktop from the KDE Control Center.

Click the plus (+) sign next to topics you want to configure. Then select the particular item you want to configure. The following sections describe some of the features you can configure from the KDE Control Center.

Changing the display

There are several ways you can change the look-and-feel of your desktop display. Under the Appearance & Themes topic (click the plus sign), you can change Background, Colors, Fonts, Icons, Launch Feedback, Panel, Screen Saver, Style, Theme Manager, and Window Decoration.

Here are a few of the desktop features you may want to change:

Change the background — Under the Appearance & Themes heading in the KDE Control Center, select Background. You can remove the checkmark next to Common Background to choose to have a common background for all four of your virtual desktops or assign backgrounds to individual desktops.

First select the Mode: Flat (single color); Pattern (two colors and click Setup to select a pattern); Background Program (click Setup to choose a program to run on the background); Horizontal Gradient (two colors that fade from left to right); Vertical Gradient (two colors that fade from top to bottom); Pyramid Gradient (two colors that fade from outside in); Pipecross Gradient (two colors that fade from the outside to a cross in the middle); and Elliptic Gradient (two colors that fade from outside to an ellipse in the middle).

If you prefer to use a wallpaper on the background, click the Wallpaper tab and select a desktop. Choose a Mode (to indicate the position of the wallpaper) and click the Wallpaper box to scroll down to choose the wallpaper. If you have a JPEG image you would like to use instead, click the Browse button to select the image you want from your file system. Click the Multiple box if you want to assign several wallpapers that change at set intervals.

Click Apply to apply your selections.

Change the screensaver — Under the Appearance & Themes heading, select Screen Saver. From the window that appears, select from about 160 different screen savers. My favorite is Slide Show, where you can have a slide show of images for your screen saver. Click Setup to identify an image directory or otherwise modify the behavior of the screen saver. Under settings, select how many minutes of inactivity before the screen saver turns on. You can also click Require Password to require that a password be entered before you can access your display after the screen saver has come on.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you are working in a place where you want your desktop to be secure, be sure to turn on the Require Password feature. This prevents others from gaining access to your comp

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Change fonts — You can assign different fonts to different places in which fonts appear on the desktop. Under the Appearance and Themes heading, select Fonts. Select one of the categories of fonts (General, Fixed width, Toolbar, Menu, Window title, Taskbar and Desktop fonts). Then click the Choose box to select a font from the Select Font box that you want to assign to that category. If the font is available, you will see an example of the text in the Sample text box.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip To use 100 dpi fonts, you need to add an entry for 100 dpi fonts to /etc/XF86Config file. After you make that change, you need to restart the X server for it to take effect.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Other attributes you can change for the selected fonts are size (in points) and character set (to select an ISO standard character set). Select Apply to apply the changes.

Change the colors — Under the Appearance & Themes heading in the KDE Control Center, select Colors. The window that appears lets you change the color of selected items on the desktop. Select a whole color scheme from the Color Scheme box. Or select an item from the Widget color box to change a particular item. Items you can change include text, backgrounds, links, buttons, and title bars.

Changing panel attributes

For most people, the Panel is the place where they select which desktop is active and which applications are run. You can change some of the behavior of the panel from the KDE Control Center. (Select Panels under the Appearance & Themes heading.) However, you can do more to modify the panel from the panel itself.

Right-click on any empty space on your panel, and then select Configure Panel. You can change these features from the Settings window that appears:

Arrangement — Change the location of the panel by clicking on Top, Left, Bottom, or Right in the Panel Location box. The Panel Style selection lets you change the size of the Panel from Medium to Tiny, Small, or Large.

Hiding — Certain selections allow you to autohide the panel or use hide buttons. Under the Hide Mode heading, choose whether to hide only when a panel hiding button is clicked or to hide automatically after a set number of seconds when the cursor is not in the panel area. You can also show or not show hiding buttons. Sliders let you select the delay and speed at which panels and buttons are hidden.

Menus — Unlike with the GNOME main menu, you have the ability to manipulate the main menu from the GUI in KDE. Click the Edit K Menu button. The KDE Menu editor that appears lets you cut, copy, paste, remove, and modify submenus and applications from your main menu.

Adding application launchers and MIME types

You want to be able to quickly access the applications that you use most often. One of the best ways to make that possible is to add icons to the panel or the desktop that can launch the applications you need with a single click. Procedures for adding applications to the panel and desktop are described in the following sections.

Adding applications to the panel

You can add any KDE application to the KDE panel quite easily. Here's how:

1. Right-click an open space on the panel.
2. Choose Add → Application Button.
3. Select one of the categories of applications.
4. Click any application from the category you selected (or select Add this menu to add the whole menu of applications).

An icon representing the application should immediately appear on the panel. (If the panel seems a bit crowded, you might want to remove some applications you don't use.) At this point, you can change any properties associated with the application by right-clicking on the application icon in the panel and then selecting Preferences. Here are the properties you can change:

General — Lets you change the filename associated with the icon. (This is not recommended, since the filename is not exposed from the panel.) You can also click the icon on this tab to select a different icon to represent it.

Permissions — Lets you change read, write, and execute permissions associated with the application. Because this is a special KDE link file (.kdelnk or .desktop), it only needs to be readable by all and not executable.

Execute — Lets you change the complete command line that is run when the command is executed, as well as the icon representing the application. Click Run

in terminal if the application is a text-based, and not a GUI-based, application.
Application — Lets you add a Comment to the application. This comment is displayed as a tool tip when the pointer is positioned over the icon. The Name you enter is what appears in the title bar for the application. If you want the application to be launched if a particular type of data (MIME Type) is encountered, you can click a File Type in the right text box, then click the left arrow button to move the MIME Type to the left box.

If you decide later that you no longer want this application to be available on the panel, right-click it and click Remove. To move it to a different location on the panel, right-click it, click Move, move it to where you want it on the panel, and click again.

Adding applications to the desktop

To add an application to the desktop, you can use the desktop menu. Here's how:

1. Right-click an open area of the desktop.
2. Select Create New → Link to Application from the menu.
3. In the Properties window that appears, replace the Program name. (This name will appear under the icon on the desktop.)
4. Add Permissions, Execute, and Application information for the application as described in the previous section about adding applications to the panel. In this case, much of the information will have to be entered manually. You will at least want to enter the command to run, along with any options, on the Execute tab. You will probably also want to enter the Comment and Name on the Application tab. Click OK to save the application icon to the desktop.

If you decide later that you no longer want this application to be available on the desktop, right-click it and click Delete or Move to Trash.

Team LiB

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Troubleshooting Your Desktop

If your desktop is not functioning properly (or at all) it may be that your video card was not configured properly. This section helps you get your video card configured properly and your desktop up and running smoothly.

GUI doesn't work at start-up

If Red Hat Linux has been successfully installed (along with the desired desktop environment) but the GUI wasn't set to start at boot time, you may only see a simple text-based login prompt when you start Red Hat Linux. This login prompt may look something like this:

Fedora Core release 1

Kernel 2.4.22 on an i686

YourComputer login:

Log in as the root user. As noted earlier, you can check if you have a GUI that is at least working well enough for you to correct it. Type the following command:

startx

If the desktop works fine when you type **`startx`**, you might want to change to a graphical login, so the GUI starts automatically every time. See the "What Happens During Desktop Startup" sidebar for information on booting to a GUI.

If your GUI is so distorted you can't even see to correct it, switch to a virtual terminal to correct the problem. For example, hold the Ctrl and Alt keys, and press F2. You will see a plain text login prompt. Log in as root user and type **init 3** to make the garbled GUI login screen go away. Then you can try tuning your video card as described in the following section.

Tip Switching virtual terminals is a great way to get out of a GUI that is broken or stuck and run the commands you need to fix a problem. You can use any function key from F1 through F8 with Ctrl+Alt to switch terminals. The GUI itself is probably on the F7 virtual terminal. Linux experts use virtual terminals during Red Hat installation to debug a problem or during startup to view text startup messages.

[redacted]

What Happens During Desktop Startup?

The X [redacted]raphical login screen is started by the `prefdm` script. By default, the login screen is disp[redacted] GNOME display manager (`gdm` command), which handles both logging in and starting the desktop environment for your console monitor, as well as graphical logins from other computers and X terminals. (If you don't like the way the login screen appears, run `gdmsetup` as root user. It will let you change the login greeter screen or even allow no-password logins.)

The `prefdm` script is launched only if the run level in the `/etc/inittab` file is set to `5`, as follows:

id:5:initdefault:

If the `initdefault` state is `3`, the system boots to a text-based login prompt. See Chapter 12 for information on Linux run states and start-up processes.

If you are unable to get the video card and monitor configured properly, or if you don't need a GUI, you can configure the computer to start up in text mode. To do this using any shell text editor (such as the `vi` command described in Chapter 3), change the `initdefault` line in the `/etc/inittab` file from `id:5:initdefault:` to `id:3:initdefault`.

If you prefer to have Red Hat Linux boot to a GUI, change the 3 to a 5 instead.

[redacted]

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tuning your video card and monitor

If your GUI is starting up but needs some tuning (to get better resolution, more colors, or to fix flickering), you can use the Display Settings window to fix your desktop. For the current Red Hat Linux, the Display Settings window was enhanced so that you can use it from a command with no GUI running. The next sections describe how to run the Display Settings window, then how to review the resulting XF86Config file to understand your settings.

Running the Display Settings window

Red Hat recently replaced the `Xconfigurator` tool with a new Display Settings window (`redhat-config-xfree86 command`). This window lets you set the most basic functions relating to your display, monitor, and video card. The Display Settings window is easy-to-use and no longer requires a running X desktop to use it.

To open the Display Settings window from the Red Hat menu, click System Settings → Display. To open that window from a text prompt (even with no GUI running), type **`redhat-config-xfree86`** (as root). The Display Settings window appears, as shown in Figure 3-20.

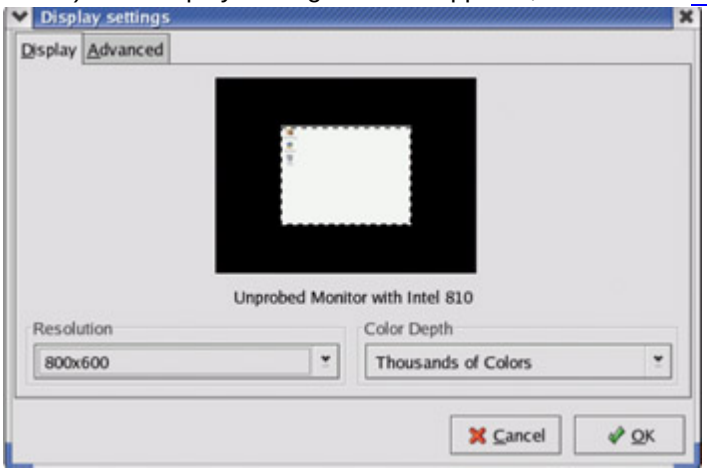

Figure 3-20: Use the Display Settings window to configure basic desktop, video card, and monitor settings.

From the Display tab, you can try different resolutions (screen width and height in pixels) and color depth (from 256 colors to millions of color). Click the Advanced tab to try to configure your monitor and video card. Click OK to save your changes.

Here are a few tips for using the Display Settings window:

- If you know your monitor type, but it is not being detected, click the Advanced tab and then click Configure. You can select the monitor from a list of monitors (by manufacturer) or, if it's not on the list, enter information about the monitor's horizontal and vertical sync rates from the manufacturer's instructions.
- If the video card supports 3D hardware acceleration, but is not turned on, you might be able to enable that feature from the Advanced tab under the Video Card section. Just click the Enable Hardware 3D Acceleration check box.

Changes made in the Display Settings window result in the creation of a new `/etc/X11/XF86Config` file. The next section describes what the `XF86Config` file contains.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If the Display Settings window fails to create a working XF86Config file, you can try another approach. With no GUI on as root user, type the following commands from a shell:

```
#
XF
re
e8
6
-co
nfi
gu
re
#
XF
re
e8
6
-xf
86
co
nfi
g
/ro
ot/
XF
86
Co
nfi
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Understanding the XF86Config file

The `XF86Config` file (located in the `/etc/X11` directory) contains definitions used by the X server to use your video card, keyboard, mouse, and monitor. In general, novice users should not edit this file directly. For some video cards, however, manual configuration may be required to get the card working properly.

The following is a description of the basic information contained in the `XF86Config` file:

ServerLayout section — Binds input and output devices for your X session. Lets you set server definitions for different X servers (if necessary).

Module section — Describes which X server modules should be loaded.

Files section — Sets the locations of the RGB (color) and fonts databases.

InputDevice sections — Separate sections identify keyboard and mouse input devices.

Monitor section — Sets the type of monitor, along with its horizontal sync rate, vertical refresh rate, and settings needed to operate at different resolutions.

Device section — Identifies your video card and, optionally, video RAM and clock information for the chipset.

Screen section — Binds the graphics board and monitor information to be referenced later by the ServerLayout section.

Keyboard section — Sets keyboard settings, including the layout of the keyboard and the way certain key sequences are mapped to the keyboard.

Pointer section — Selects the pointer you are using (typically a mouse linked to `/dev/mouse`). Also sets speed and button emulation, when appropriate.

DRI — Provides information for Direct Rendering Infrastructure (used for accelerated 3D graphics).

Configuring video cards for gaming

Some games and video players require special features to work properly (or at all, in some cases). For games that require 3D hardware acceleration, including some that run under TransGaming's WineX, TransGaming recommends using nVIDIA GeForce Graphics cards.

Because only basic nVIDIA video card drivers are included in Red Hat Linux (nVIDIA's own drivers are not open source), you need to get nVIDIA drivers yourself to use those cards for gaming. You can download nVIDIA drivers for Linux from:

www.nvidia.com/object/linux_display_archive.html

Select the latest drivers from the list (look for version 1.0-4496 or later). Download and install the drivers as described.

Games that don't require 3D hardware acceleration should work fine with most video cards that are supported by XFree86 4.0 and higher drivers.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip To use hardware DRI acceleration on Voodoo 3 cards, you must have your display set to use 16 bpp resolution. On Voodoo 5 cards, only 16 bpp and 24 bpp resolutions are sup

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Getting more information

If you tried configuring X and you still have a server that crashes or has a garbled display, your video card may either be unsupported or may require special configuration. Here are several locations you can check for further information:

XFree86.org (`www.xfree86.org`) — The latest information about the X servers that come with Red Hat Linux is available from the XFree86.org Web site. XFree86 is the freeware version of X used with all major Linux distributions.

Red Hat Support (`www.redhat.com/support`) — Search the Red Hat support database for the model of your card. There may already be reports of problems (and hopefully fixes) related to your card.

X documentation — README files that are specific to different types of video cards are delivered with XFree86. Visit the XFree86 `doc` directory (`/usr/X11R6/lib/X11/doc`) for a README file specific to the type of video card (or more specifically, the video chipset) you are using. A lot of good information can also be found on the XF86Config man page (type **man XF86Config**).

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

The X Window System provides the basis for most graphical user interfaces available for Red Hat and other Linux systems today. Although X provides the framework for running and sharing applications, the GNOME and KDE desktop environments, along with a window manager and theme, provide the look and feel of your desktop.

Using various configuration files and commands, you can change nearly every aspect of your graphical environment. Backgrounds can be assigned a single color or can be filled with single or tiled graphic images. Menus can be changed or enhanced. Multiple virtual workspaces can be used and managed. After reading this chapter, you should feel comfortable working with the GNOME and KDE desktops. The next chapter should help you work from the traditional command line interface, referred to as the shell.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 4: Using Linux Commands

Overview

This chapter presents a view of Red Hat Linux from the shell. The *shell* is a command-line interpreter that lets you access some of the most critical Red Hat Linux tools. The shell is powerful, complex, and almost completely unintuitive.

Although, at first, it isn't obvious how to use the shell, with the right help you can quickly learn many of the most important shell features. This chapter is your guide to working with the Linux system commands, processes, and file system from the shell. It describes the shell environment and helps you tailor it to your needs. It also describes how to use and move around the file system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The Shell Interface

There are several ways to get to a shell interface in Red Hat Linux. Common ways are:

No GUI — If your Red Hat Linux system has no GUI (or one that isn't working at the moment), you must enter commands from the shell.

A Terminal window — With the desktop GUI running, you can open a Terminal window (right-click on the desktop, then click New Terminal) to start a shell. You can begin typing commands into the Terminal window.

If you are using a shell interface, the first thing you see is the shell prompt. The default prompt for a user is simply a dollar sign:

$

The default prompt for the root user is a pound sign (also called a hash mark):

#

For most Red Hat Linux systems, the $ or # prompts are preceded by your user name, system name, and current directory name. So, for example, a login prompt for the user named `jake` on a computer named `pine` with `/tmp` as the current directory would appear as:

[jake@pine tmp]$

You can change the prompt to display any characters you like. You could use as your prompt the current directory, the date, the local computer name, or any string of characters. (To configure your prompt, see the "Setting your prompt" section later in this chapter.)

Although there are a tremendous number of features available with the shell, it's easy to begin by just typing a few commands. Try some of the commands shown in the remainder of this section to become familiar with your current shell environment.

Tip If, instead of a shell prompt, you see a GUI when you log in, you can still try out the shell commands shown in the next section. To access a shell from the GUI, you can open a Terminal window by right-clicking on your desktop and selecting New Terminal.

In the examples that follow, the $ or # symbols indicate a prompt. The prompt is followed by the command that you type and then by Enter or Return (depending on your keyboard). The lines that follow show the output that results from the command.

Checking your login session

When you log in to a Linux system, Linux views you as having a particular identity. That identity includes your user name, group name, user ID, and group ID. Linux also keeps track of your login session: it knows when you logged in, how long you have been idle, and where you logged in from. To find out information about your identity, use the `id` command as follows:

$ **id**

uid=501(chris) gid=105(sales) groups=105(sales),4(adm),7(lp)

This shows that the user name is `chris`, which is represented by the numeric user ID (uid) `501`. Here, the primary group for chris is called `sales`, which has a group ID (gid) of `105`. Chris also belongs to other groups called `adm` (uid 4) and `lp` (uid 7). These names and numbers represent the permissions that chris has to access computer resources. (Permissions are described later in this chapter in the section on working with files.)

You can see information about your current login session by using the `who` command. In the following example, the `-m` option tells the `who` command to print information about the current user, `-u` says to add information about idle time and the process ID, and `-H` asks that a header be printed:

$ **who -umH**

NAME LINE TIME IDLE PID COMMENT

chris tty1 Jan 13 20:57 . 2013

The output from this `who` command shows that the user name is `chris`. Here, `chris` is logged in on `tty1` (which is the monitor connected to the computer), and his login session began at 20:57 on

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

January 13. The IDLE time shows how long the shell has been open without any command being typed (the dot indicates that it is currently active). COMMENT would show the name of the remote computer the user had logged in from, if that user had logged in from another computer on the network, or the name of the local X display if you were using a Terminal window (such as :0.0).

Checking directories and permissions

Associated with each shell is a location in the Linux file system known as the *current* or *working directory.* As previously mentioned, each user has a directory that is identified as the user's home directory. When you first log in to Linux, you begin with your home directory as the current directory.

When you request to open or save a file, your shell uses the current directory as the point of reference. Simply give a filename when you save a file, and it will be placed in the current directory. Alternatively, you can identify a file by its relation to the current directory (relative path). Or you can ignore the current directory and identify a file by the full directory hierarchy that locates it (absolute path). The structure and use of the file system is described in detail later in this chapter.

To find out what your current directory is, type the pwd command:

```
$ pwd
/usr/bin
```

In this example, the current/working directory is /usr/bin. To find out the name of your home directory, type the echo command, followed by the $HOME variable:

```
$ echo $HOME
/home/chris
```

In the preceding example, the home directory is /home/chris. To get back to your home directory, you can simply type the change directory (cd) command. Although cd, followed by a directory name, changes the current directory to the directory that you choose, simply typing cd takes you to your home directory:

```
$ cd
```

At this point, list the contents of your home directory, using the ls command. Either you can type the full path to your home directory to list its contents, or you can use the ls command without a directory name to list the contents of the current directory. Using the -a option to ls enables you to view the hidden files (dot files) as well as all other files. With the -l option, you can see a long, detailed list of information on each file. (You can put multiple single-letter options together after a single dash, for example, -la.)

```
$ ls -la /home/chris
total 158
drwxrwxrwx    2   chris   sales    1024  May 12 13:55 .
drwxr-xr-x    3   root    root     1024  May 10 01:49 ..
-rw-------    1   chris   sales    2204  May 18 21:30 .bash_history
-rw-r--r--    1   chris   sales      24  May 10 01:50 .bash_logout
-rw-r--r--    1   chris   sales     230  May 10 01:50 .bash_profile
-rw-r--r--    1   chris   sales     124  May 10 01:50 .bashrc
drw-r--r--    1   chris   sales    4096  May 10 01:50 .kde
-rw-rw-r--    1   chris   sales  149872  May 11 22:49 letter
```

Displaying a long list (-l option) of the contents of your home directory shows you more about file sizes and directories. Directories such as the current directory (.) and the directory above the current directory (..) are noted as directories by the letter "d" at the beginning of each entry. In this case, dot (.) represents /home/chris and two dots (..), which is also referred to as the parent directory, represents /home. The /home directory is owned by root. All other files are owned by the user chris (who belongs to the sales group).

The file or directory names shown on the right are mostly dot (.) files that are used to store GUI properties (.kde) or shell properties (.bash files). The only non-dot file shown in this example is the one named letter. At the beginning of each line are the permissions set for each file. (Permissions and configuring shell property files are described later in this chapter.) Other information in the listing includes the size of each file in bytes (column 4) and the date and time each file was most recently

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

modified (column 5).

Checking system activity

In addition to being a multiuser operating system, Linux is also a multitasking system. *Multitasking* means that many programs can be running at the same time. An instance of a running program is referred to as a *process*. Linux provides tools for listing running processes, monitoring system usage, and stopping (or killing) processes when necessary.

The most common utility for checking running processes is the `ps` command. With `ps`, you can see which programs are running, the resources they are using, and who is running them. The following is an example of the `ps` command:

```
$ ps -au
USER   PID %CPU %MEM  VSZ    RSS   TTY    STAT START  TIME COMMAND
root   2146 0.0  0.8 1908   1100   ttyp0  S    14:50  0:00 login -- jake
jake   2147 0.0  0.7 1836   1020   ttyp0  S    14:50  0:00 -bash
jake   2310 0.0  0.7 2592    912   ttyp0  R    18:22  0:00 ps -au
```

In this example, the `-a` option asks to show processes of all users that are associated with your current terminal, and the `-u` option asks that user names be shown, as well as other information such as the time the process started and memory and CPU usage. The concept of terminal comes from the old days, when people worked exclusively from character terminals, so a terminal typically represented a single person at a single screen. Now you can have many "terminals" on one screen by opening multiple Terminal windows.

On this shell session, there isn't much happening. The first process shows that the user named jake logged in to the login process (which is controlled by the root user). The next process shows that jake is using a bash shell and has just run the `ps au` command. The terminal device `ttyp0` is being used for the login session. The `STAT` column represents the state of the process, with `R` indicating a currently running process and `S` representing a sleeping process.

The `USER` column shows the name of the user who started the process. Each process is represented by a unique ID number referred to as a process ID (PID). (You can use the PID if you ever need to kill a runaway process.) The `%CPU` and `%MEM` columns show the percentage of the processor and random access memory, respectively, that the process is consuming. `VSZ` (virtual set size) shows size of the image process (in kilobytes), and `RSS` (resident set size) shows the size of the program in memory. `START` shows the time the process began running, and `TIME` shows the cumulative system time used.

Many processes running on a computer are not associated with a terminal. A normal Red Hat Linux system has many processes running in the background. Background system processes perform such tasks as logging system activity or listening for data coming in from the network. They are often started when Red Hat Linux boots up and runs continuously until it shuts down. To see and thereby monitor all the processes running on your Red Hat Linux system, type:

$ **ps aux | less**

I added the pipe (`|`) and the `less` command to `ps aux` to allow you to page through the many processes that will appear on your screen. Use the spacebar to page through and type **q** to end the list. A pipe lets you direct the output of one command to be the input of the next command.

Exiting the shell

To exit the shell when you are done, either type **exit** or press Ctrl+D. If you are exiting from your login shell (the shell that started when you first logged in), type **logout** to exit the shell.

I just showed a few commands designed to familiarize you quickly with your Linux system. There are hundreds of other commands that you can try that are contained in directories such as `/bin` and `/usr/bin`. There are also administrative commands in `/sbin` or `/usr/sbin` directories. Many of these commands are described in the remainder of this chapter.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Understanding the Red Hat Linux Shell

Before icons and windows took over computer screens, you typed commands to run most computers. On UNIX systems, from which Red Hat Linux was derived, the program used to interpret and manage commands was referred to as the shell.

The shell provides a way to run programs, work with the file system, compile computer code, and manage the computer. Although the shell is less intuitive than common GUIs, most Linux experts consider the shell to be much more powerful than GUIs. Because shells have been around for so long, many advanced features have been built into them. Many old-school Linux administrators and programmers primarily use a GUI as a way to open lots of shells.

The Red Hat Linux shell illustrated in this chapter is called the *bash* shell, which stands for Bourne Again SHell. The name is derived from the fact that bash is compatible with the first UNIX shell: the Bourne shell (represented by the `sh` command). Other popular shells include the C Shell (*csh*), which is popular among BSD UNIX users, and the Korn Shell (*ksh*), which is popular among UNIX System V users. Linux also has a *tcsh* shell (a C shell look-alike) and an *ash* shell (another Bourne shell look-alike).

Although most Red Hat Linux users have a preference for one shell or another, when you know how to use one shell, you can quickly learn any of the others by occasionally referring to the shell's man page (for example, type `man bash`). In Red Hat Linux, the bash shell is roughly compatible with the `sh` shell.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the Shell in Red Hat Linux

When you type a command in a shell, you can also include other characters that change or add to how the command works. In addition to the command itself, these are some of the other items that you can type on a shell command line:

Options — Most commands have one or more options you can add to change their behavior. Options typically consist of a single letter, preceded by a dash. You can also often combine several options after a single dash. For example, the command `ls -la` lists the contents of the current directory. The `-l` asks for a detailed (long) list of information, and the `-a` asks that files beginning with a dot (.) also be listed. When a single option consists of a word, it is usually preceded by a double dash (`--`). For example, to use the help option on many commands, you would enter `--help` on the command line.

Arguments — Many commands also accept arguments after any options are entered. An argument is an extra piece of information, such as a filename, that can be used by the command. For example, `cat /etc/passwd` displays the contents of the `/etc/passwd` file on your screen. In this case, `/etc/passwd` is the argument.

Environment variables — The shell itself stores information that may be useful to the user's shell session in what are called *environment variables.* Examples of environment variables include `$SHELL` (which identifies the shell you are using), `$PS1` (which defines your shell prompt), and `$MAIL` (which identifies the location of your mailbox).

Tip You can check your environment variables at any time. Type `declare` to list the current environment variables. Or you can type `echo $`*VALUE*, where *VALUE* is replaced by the name of a particular environment variable you want to list.

Metacharacters — These are characters that have special meaning to the shell. Metacharacters can be used to direct the output of a command to a file (`>`), pipe the output to another command (`|`), or run a command in the background (`&`), to name a few. Metacharacters are discussed later in this chapter.

To save you some typing, there are shell features that store commands you want to reuse, recall previous commands, and edit commands. You can create aliases that allow you to type a short command to run a longer one. The shell stores previously entered commands in a history list, which you can display and from which you can recall commands. This is discussed further in the remainder of this section.

Unless you specifically change to another shell, the bash shell is the one you use with Red Hat Linux. The bash shell contains most of the powerful features available in other shells. Although the description in this chapter steps you through many bash shell features, you can learn more about the bash shell by typing `man bash`. For other ways to learn about using the shell, refer to the sidebar "Getting Help with Using the Shell."

Locating commands

If you know the directory that contains the command you want to run, one way to run it is to type the full path to that command. For example, you run the `date` command from the `/bin` directory by typing:

$ **/bin/date**

Of course, this can be inconvenient, especially if the command resides in a directory with a long name. The better way is to have commands stored in well-known directories, and then add those directories to your shell's PATH environment variable. The path consists of a list of directories that are checked sequentially for the commands you enter. To see your current path, type the following:

$ echo $PATH

/bin:/usr/bin:/usr/local/bin:/usr/bin/X11:/usr/X11R6/bin:/home/chris/bin

The results show the default path for a regular Linux user. Directories in the path list are separated by colons. Most user commands that come with Linux are stored in the `/bin`, `/usr/bin`, or `/usr/local/bin` directories. Graphical commands (that are used with GUIs) are contained in `/usr/bin/X11` and `/usr/X11R6/bin` directories. The last directory shown is the bin directory in the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

user's home directory.

Tip If you want to add your own commands or shell scripts, place them in the bin directory in your home directory (such as /home/chris/bin for the user named chris). This directory is automatically added to your path. So as long as you add the command to your bin with execute permission (described in the "Understanding file permissions" section), you can immediately begin using the command by simply typing the command name at your shell prompt.

Getting Help with Using the Shell

When [redacted]t using the shell, it can be intimidating. All you see is a prompt. How do you know which commands are available, which options they use, or how to use more advanced features? Fortunately, lots of help is available. Here are some places you can look to supplement what you learn in this chapter:

Check the PATH — Type **echo $PATH**. You see a list of the directories containing commands that are immediately accessible to you. Listing the contents of those directories displays most standard Linux commands.
Use the `help` command — Some commands are built into the shell, so they do not appear in a directory. The `help` command lists those commands and shows options available with each of them. (Type `help | less` to page through the list.) For help with a particular built-in command, type `help` *`command`*, replacing *`command`* with the name that interests you.
Use `--help` with the command — Many commands include a `--help` option that you can use to get information about how the command is used. For example, type `date --help | less`. The output shows not only options, but also time formats you can use with the date command.
Use the `man` command — To learn more about a particular command, type `man` *`command`*. (Replace *command* with the command name you want.) A description of the command and its options appears on the screen.

If you are the root user, directories containing administrative commands are in your path. These directories include `/sbin` and `/usr/sbin`.
The path directory order is important. Directories are checked from left to right. So, in this example, if there is a command called `foo` located in both the `/bin` and `/usr/bin` directories, the one in `/bin` is executed. To have the other `foo` command run, you either type the full path to the command or change your `PATH` variable. (Changing your `PATH` and adding directories to it are described later in this chapter.)

Not all the commands that you run are located in directories in your PATH. Some commands are built into the shell. Other commands can be overridden by creating aliases that define any commands and options that you want the command to run. There are also ways of defining a function that consists of a stored series of commands. Here is the order in which the shell checks for the commands you type:

1. **Aliases —** Names set by the `alias` command that represent a particular command and a set of options. (Type `alias` to see what aliases are set.)
2. **Shell reserved word —** Words that are reserved by the shell for special use. Many of these are words that you would use in programming-type functions, such as `do`, `while`, `case`, and `else`.
3. **Function —** A set of commands that are executed together within the current shell.
4. **Built-in command —** A command that is built into the shell.
5. **File system command —** This is a command that is stored in and executed from the computer's file system. (These are the commands that are indicated by the value

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

of the `PATH` variable.)

To find out where a particular command is taken from, you can use the `type` command. For example, to find out where the bash shell command is located, type the following:

```
$ type bash
bash is /bin/bash
```

Try these few words with the `type` command to see other locations of commands: `which`, `case`, and `return`. If a command resides in several locations, you can add the `-a` option to have all the known locations of the command printed.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Sometimes you run a command and receive an error message that the command was not found or that permission to run the command was denied. In the first case, check that

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Rerunning commands

It's annoying, after typing a long or complex command line, to learn that you mistyped something. Fortunately, some shell features let you recall previous command lines, edit those lines, or complete a partially typed command line.
The *shell history* is a list of the commands that you have entered before. Using the `history` command, you can view your previous commands. Then, using various shell features, you can recall individual command lines from that list and change them however you please.

The rest of this section describes how to do command-line editing, how to complete parts of command lines, and how to recall and work with the history list.

Command-line editing

If you type something wrong on a command line, the bash shell ensures that you don't have to delete the entire line and start over. Likewise, you can recall a previous command line and change the elements to make a new command.
By default, the bash shell uses command-line editing that is based on the `emacs` text editor. So, if you are familiar with `emacs`, you probably already know most of the keystrokes described here.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you prefer the vi command for editing shell command lines, you can easily make that happen. Add the line set -o vi to the .bashrc file in your home directory. The next time you open a shell, you can use vi commands (as described in the tutorial later in this chapter) to edit your command lines.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To do the editing, you can use a combination of control keys, meta keys, and arrow keys. For example, Ctrl+f means to hold the control key and type f. Alt+f means to hold the Alt key and type f. (Instead of the Alt key, your keyboard may use a Meta key or the Esc key instead. On a Windows keyboard, use the Windows key.)

To try out a bit of command-line editing, type the following command:

$ **ls /usr/bin | sort -f | more**
This command lists the contents of the `/usr/bin` directory, sorts the contents in alphabetical order (regardless of upper- and lowercase), and pipes the output to `more` (so you can page through the results). Now, suppose you want to change `/usr/bin` to `/bin`. You can use the following steps to change the command:

1. Press Ctrl+a. This moves the cursor to the beginning of the command line.
2. Press Ctrl+f or the right arrow (→) key. Repeat this command a few times to position the cursor under the first slash (/).
3. Press Ctrl+d. Type this command four times to delete `/usr`.
4. Press Enter. This executes the command line.

As you edit a command line, at any point you can type regular characters to add those characters to the command line. The characters appear at the location of your cursor. You can use right (→) and left (() arrows to move the cursor from one end to the other on the command line. You can also press the up (() and down (() arrow keys to step through previous commands in the history list to select a command line for editing. (See the discussion on command recall for details on how to recall commands from the history list.)
There are many keystrokes you can use to edit your command lines. Table 4-1 lists the keystrokes that you can use to move around the command line.

Table 4-1: Keystrokes for Navigating Command Lines

Keystroke	Full Name	Meaning
Ctrl+f	Character forward.	Go forward one character.
Ctrl+b	Character backward.	Go backward one character.
Alt+f	Word forward.	Go forward one word.
Alt+b	Word backward.	Go backward one word.
Ctrl+a	Beginning of line.	Go to the beginning of the current line.
Ctrl+e	End of line.	Go to the end of the line.
Ctrl+l	Clear screen.	Clear screen and leave line at the top of the screen.

Table 4-2 lists the keystrokes for editing command lines.

Table 4-2: Keystrokes for Editing Command Lines

Keystroke	Full Name	Meaning
Ctrl+d	Delete current.	Delete the current character.
Backspace or Rubout	Delete previous.	Delete the previous character.
Ctrl+t	Transpose character.	Switch positions of current and previous characters.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 4-2: Keystrokes for Editing Command Lines

Keystroke	Full Name	Meaning
Alt+t	Transpose words.	Switch positions of current and previous characters.
Alt+u	Uppercase word.	Change the current word to uppercase.
Alt+l	Lowercase word.	Change the current word to lowercase.
Alt+c	Capitalize word.	Change the current word to an initial capital.
Ctrl+v	Insert special character.	Add a special character. For example, to add a Tab character, press Ctrl+v+Tab.

Table 4-3 lists the keystrokes for cutting and pasting text on a command line.

Table 4-3: Keystrokes for Cutting and Pasting Text in Command Lines

Keystroke	Full Name	Meaning
Ctrl+k	Cut end of line.	Cut text to the end of the line.
Ctrl+u	Cut beginning of line.	Cut text to the beginning of the line.
Ctrl+w	Cut previous word.	Cut the word located behind the cursor.
Alt+d	Cut next word.	Cut the word following the cursor.
Ctrl+y	Paste recent text.	Paste most recently cut text.
Alt+y	Paste earlier text.	Rotate back to previously cut text and paste it.
Ctrl+c	Delete whole line.	Delete the entire line.

Command line completion

To save you a few keystrokes, the bash shell offers several different ways of completing partially typed values. To attempt to complete a value, type the first few characters, and then press Tab. Here are some of the values you can type partially:

Environment variable — If the text begins with a dollar sign ($), the shell completes the text with an environment variable from the current shell.
User name — If the text begins with a tilde (~), the shell completes the text with a user name.
Command, alias, or function — If the text begins with regular characters, the shell tries to complete the text with a command, alias, or function name.
Host name — If the text begins with an at (@) sign, the shell completes the text with a host name taken from the `/etc/hosts` file.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip To add hostnames from an additional file, you can set the HOSTFILE variable to the name of that file. The file must be in the same format as /etc/hosts.

Here are a few examples of command completion. (When you see <Tab>, it means to press the Tab key on your keyboard.) Type the following:

$ ***echo*** **$OS<Tab>**

$ **cd ~ro<*Tab*>**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

$ **fing<*Tab*>**

$ **mailx root@loc<*Tab*>**

The first example causes `$OS` to expand to the `$OSTYPE` variable. In the next example, `~ro` expands to the root user's home directory (`~root/`). Next, `fing` expands to the `finger` command. Finally, the address of `root@loc` expands to computer name `localhost`.

Of course, there will be times when there are several possible completions for the string of characters you have entered. In that case, you can check the possible ways text can be expanded by pressing Esc+? (or by pressing Tab twice) at the point where you want to do completion. This shows the result you would get if you checked for possible completions on `$P`.

$ **echo $P<Esc+?>**

$PATH $PPID $PS1 $PS2 $PS4 $PWD

$ echo $P

In this case, there are six possible variables that begin with `$P`. After possibilities are displayed, the original command line returns, ready for you to complete it as you choose.

If text you are trying to complete is not preceded by a `$`, `~`, or `@`, you can still try to complete the text with a variable, user name, or host name. Press the following to complete your text:

Alt+~ — Complete the text before this point as a user name.
Alt+$ — Complete the text before this point as a variable.
Alt+@ — Complete the text before this point as a host name.
Alt+! — Complete the text before this point as a command name (alias, reserved word, shell function, built-in, and filenames are checked in that order). In other words, complete this key sequence with a command that you previously ran.
Ctrl+x+/ — List possible user name text completions.
Ctrl+x+$ — List possible environment variable completions.
Ctrl+x+@ — List possible host name completions.
Ctrl+x+! — List possible command name completions.

Command line recall

After you type a command line, that entire command line is saved in your shell's history list. The list is stored in a history file, from which any command can be recalled to run again. After it is recalled, you can modify the command line, as described earlier.

To view your history list, use the `history` command. Type the command without options or followed by a number to list that many of the most recent commands. For example:

$ **history 8**

```
 382 date
 383 ls /usr/bin | sort -a | more
 384 man sort
 385 cd /usr/local/bin
 386 man more
 387 useradd -m /home/chris -u 101 chris
 388 passwd chris
 389 history 8
```

A number precedes each command line in the list. There are several ways to run a command immediately from this list, including:

Run Command Number (`!`*n*) — Replace the *n* with the number of the command line, and the command line indicated is run.
Run Previous Command (`!!`) — Runs the previous command line.
Run Command Containing String (`!?`*string*`?`) — Runs the most recent command that contains a particular *string* of characters.

Instead of just running a `history` command line immediately, you can recall a particular line and edit it. You can use these keys to do that:

Step (Arrow Keys) — Press the up (() and down (() arrow keys to step through

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

each command line in your history list to arrive at the one you want. (Ctrl+p and Ctrl+n do the same functions, respectively.)

Reverse Incremental Search (Ctrl+r) — After you press these keys, you are asked to enter a search string to do a reverse search. As you type the string, a matching command line appears that you can run or edit.

Forward Incremental Search (Ctrl+s) — After you press these keys, you are asked to enter a search string to do a forward search. As you type the string, a matching command line appears that you can run or edit.

Reverse Search (Alt+p) — After you press these keys, you are asked to enter a string to do a reverse search. Type a string and press Enter to see the most recent command line that includes that string.

Forward Search (Alt+n) — After you press these keys, you are asked to enter a string to do a forward search. Type a string, and press Enter to see the most recent command line that includes that string.

Beginning of History List (Alt+<) — Brings you to the first entry of the history list.

Beginning of History List (Alt+>) — Brings you to the last entry of the history list.

Another way to work with your history list is to use the `fc` command. Type **fc** followed by a history line number, and that command line is opened in a text editor. Make the changes that you want. When you exit the editor, the command runs. You could also give a range of line numbers (for example, `fc 100 105`). All the commands open in your text editor, and then run one after the other when you exit the editor.

The history list is stored in the `.bash_history` file in your home directory. Up to 1000 history commands are stored for you by default.

Connecting and expanding commands

A truly powerful feature of the shell is the capability to redirect the input and output of commands to and from other commands and files. To allow commands to be strung together, the shell uses metacharacters. As noted earlier, a metacharacter is a typed character that has special meaning to the shell for connecting commands or requesting expansion.

Piping commands

The pipe (`|`) metacharacter connects the output from one command to the input of another command. This lets you have one command work on some data, then have the next command deal with the results. Here is an example of a command line that includes pipes:

$ **cat /etc/password | sort | more**

This command lists the contents of the `/etc/password` file and pipes the output to the `sort` command. The `sort` command takes the user names that begin each line of the `/etc/password` file, sorts them alphabetically, and pipes the output to the `more` command. The `more` command displays the output one page at a time, so that you can go through the output a line or a page at a time.

Pipes are an excellent illustration of how UNIX, the predecessor of Linux, was created as an operating system made up of building blocks. A standard practice in UNIX was to connect utilities in different ways to get different jobs done. For example, before the days of graphical word processors, users created plain-text files that included macros to indicate formatting. To see how the document really appeared, they would use a command such as the following:

$ gunzip < /usr/share/man/man1/grep.1.gz | nroff -c -man | less

In this example, the contents of the `grep` man page (`grep.1.gz`) are directed to the `gunzip` command to be unzipped. The output from `gunzip` is piped to the `nroff` command to format the man page using the manual macro (`-man`). The output is piped to the `less` command to display the output. Because the file being displayed is in plain text, you could have substituted any number of options to work with the text before displaying it. You could sort the contents, change or delete some of the content, or bring in text from other documents. The key is that, instead of all those features being in one program, you get results from piping and redirecting input and output between multiple commands.

Sequential commands

Sometimes you may want a sequence of commands to run, with one command completing before the next command begins. You can do this by typing several commands on the same command line and

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

separating them with semicolons (`;`):

$ date ; troff -me verylargedocument | lpr ; date
In this example, I was formatting a huge document and wanted to know how long it would take. The first command (`date`) showed the date and time before the formatting started. The `troff` command formatted the document and then piped the output to the printer. When the formatting was done, the date and time was printed again (so I knew how long the `troff` command took to complete).

Background commands

Some commands can take a while to complete. Sometimes you may not want to tie up your shell waiting for a command to finish. In those cases, you can have the commands run in the background by using the ampersand (`&`).
Text formatting commands (such as `nroff` and `troff`, described earlier) are examples of commands that are often run in the background to format a large document. You also may want to create your own shell scripts that run in the background to check continuously for certain events to occur, such as the hard disk filling up or particular users logging in.

Here is an example of a command being run in the background:

$ **troff -me verylargedocument | lpr &**
There are other ways of managing background and foreground processes (described in the "Managing background and foreground processes" section).

Expanding commands

With command substitution, you can have the output of a command interpreted by the shell instead of by the command itself. In this way, you can have the standard output of a command become an argument for another command. The two forms of command substitution are `$(`*`command`*`)` or `'`*`command`*`'`.

The command in this case can include options, metacharacters, and arguments. Here is an example of using command substitution:

$ **vi $(find /home | grep xyzzy)**
In this command line, the command substitution is done before the `vi` command is run. First, the `find` command starts at the `/home` directory and prints out all files and directories below that point in the file system. This output is piped to the `grep` command, which filters out all files except for those that include the string `xyzzy`. Finally, the `vi` command opens all filenames for editing (one at a time) that include `xyzzy`.

This particular example may be useful if you knew that you wanted to edit a file for which you knew the name but not the location. As long as the string was uncommon, you could find and open every instance of a filename existing beneath a point you choose in the file system.

Expanding arithmetic expressions

There may be times when you want to pass arithmetic results to a command. There are two forms you can use to expand an arithmetic expression and pass it to the shell: `$[`*`expression`*`]` or `$((`*`expression`*`))`. Here is an example:

$ **echo "I am $[2003 - 1957] years old."**

I am 46 years old.
In this example, the shell interprets the arithmetic expression first (`2003 – 1957`), and then passes that information to the `echo` command. The `echo` command displays the text, with the results of the arithmetic (`46`) inserted.

Expanding variables

Environment variables that store information within the shell can be expanded using the dollar sign (`$`) metacharacter. When you expand an environment variable on a command line, the value of the variable is printed instead of the variable name itself, as follows:

$ **ls -l $BASH**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

-rwxr-xr-x 1 root root 625516 Dec 5 11:13 /bin/bash
Using `$BASH` as an argument to `ls -l` causes a long listing of the bash command to be printed. For more information on shell environment variables, see the following section.

Using shell environment variables
Every active shell stores pieces of information that it needs to use in what are called *environment variables*. An environment variable can store things such as locations of configuration files, mailboxes, and path directories. They can also store values for your shell prompts, the size of your history list, and type of operating system.
To see the environment variables currently assigned to your shell, type the `declare` command. (It will probably fill more than one screen, so type `declare | more`.) You can refer to the value of any of those variables by preceding it with a dollar sign (`$`) and placing it anywhere on a command line. For example:

$ **echo $USER**

chris
This command prints the value of the `USER` variable, which holds your user name (`chris`). Substitute any other value for `USER` to print its value instead.

Common shell environment variables

When you start a shell (by logging in or opening a Terminal window), a lot of environment variables are already set. Here are some variables that are either set when you use a bash shell or that can be set by you to use with different features.

`BASH` — Contains the full path name of the `bash` command. This is usually `/bin/bash`.
`BASH_VERSION` — A number of the current version of the `bash` command.
`EUID` — This is the effective user ID number of the current user. It is assigned when the shell starts, based on the user's entry in the `/etc/passwd` file.
`FCEDIT` — If set, this indicates the text editor used by the `fc` command to edit `history` commands. If this variable isn't set, the `vi` command is used.
`HISTFILE` — The location of your history file. It is typically located at `$HOME/.bash_history`.
`HISTFILESIZE` — The number of history entries that can be stored. After this number is reached, the oldest commands are discarded. The default value is 1000.
`HISTCMD` — This returns the number of the current command in the history list.
`HOME` — This is your home directory. It is your current working directory each time you log in or type the `cd` command with any options.
`HOSTTYPE` — A value that describes the computer architecture on which the Linux system is running. For Intel-compatible PCs, the value is i386, i486, i586, or i686.
`MAIL` — This is the location of your mailbox file. The file is typically your user name in the `/var/spool/mail directory`.
`OLDPWD` — The directory that was the working directory before you changed to the current working directory.
`OSTYPE` — A name identifying the current operating system. In our case, this will say linux-gnu. (Bash can run on other operating systems as well.)
`PATH` — The colon-separated list of directories used to find commands that you type. The default value for regular users is:

/bin:/usr/bin:/usr/local/bin:/usr/bin/X11:/usr/X11R6/bin:~/bin
For the root user, the value also includes `/sbin`, `/usr/sbin`, and `/usr/local/sbin`.
`PPID` — The process ID of the command that started the current shell (for example, its parent process).
`PROMPT_COMMAND` — Can be set to a command name that is run each time before your shell prompt is displayed. Setting `PROMPT_COMMAND=date` lists the current date/time before the prompt appears.
`PS1` — Sets the value of your shell prompt. There are many items that you can read into your prompt (date, time, user name, host name, and so on). Sometimes a command requires additional prompts, which you can set with the variables `PS2`,

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

PS3, and so on. (Setting your prompt is described later in this chapter.)
PWD — This is the directory that is assigned as your current directory. This value changes each time you change directories using the cd command.
RANDOM — Accessing this variable causes a random number to be generated. The number is between 0 and 99999.
SECONDS — The number of seconds since the time the shell was started.
SHLVL — The number of shell levels associated with the current shell session. When you log in to the shell, the SHLVL is 1. Each time you start a new bash command (by, for example, using su to become a new user, or by simply typing bash), this number is incremented.
TMOUT — Can be set to a number representing the number of seconds the shell can be idle without receiving input. After the number of seconds is reached, the shell exits. This is a security feature that makes it less likely for unattended shells to be accessed by unauthorized people. (This must be set in the login shell for it to actually cause the shell to log out the user.)
UID — The user ID number assigned to your user name. The user ID number is stored in the /etc/password file.

Set your own environment variables

Environment variables can provide a handy way of storing bits of information that you use often from the shell. You can create any variables that you want (avoiding those that are already in use) so that you can read in the values of those variables as you use the shell. (The bash man page lists variables already in use.)

To set an environment variable temporarily, you can simply type a variable name and assign it to a value. Here is an example:

$ **AB=/usr/dog/contagious/ringbearer/grind ; export AB**

This example causes a long directory path to be assigned to the AB variable. The export AB command says to export the value to the shell so that it can be propagated to other shells you may open. With AB set, you can go to the directory by typing the following:

$ **cd $AB**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip You may have noticed that environment variables shown here are in all caps. Though case does matter with these variables, setting them as upper case is a convention, not a necessity

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The problem with setting environment variables in this way is that as soon as you exit the shell in which you set the variable, the setting is lost. To set variables more permanently, you should add variable settings to a bash configuration file, as described later in this section.
If you want to have other text right up against the output from an environment variable, you can surround the variable in braces. This protects the variable name from being misunderstood. For example, if you wanted to add a command name to the `AB` variable shown earlier, you could type the following:

$ **echo ${AB}/adventure**

/usr/dog/contagious/ringbearer/grind/adventure
Remember that you must export the variable so that it can be picked up by other shell commands. You must add the export line to a shell configuration file for it to take effect the next time you log in. The `export` command is fairly flexible. Instead of running the `export` command after you set the variable, you could do it all in one step, as follows:

$ **export XYZ=/home/xyz/bin**
You can override the value of any environment variable. This can be temporary by simply typing the new value. Or you can add the new export line to your `$HOME/.bashrc` file. One useful variable to update is `PATH`. Here is an example:

$ **export PATH=$PATH:/home/xyz/bin**
In this example, I added the `/home/xyz/bin` directory to the `PATH`, a useful technique if you want to run a bunch of commands from a directory that is not normally in your `PATH`, without typing the full or relative path each time.
If you decide that you no longer want a variable to be set, you can use the `unset` command to erase its value. For example, you could type `unset XYZ`, which would cause `XYZ` to have no value set. (Remember to remove the export from the `$HOME/.bashrc` file — if you added it there — or it will return the next time you open a shell.)

Managing background and foreground processes
If you are using Linux over a network or from a *dumb* terminal (a monitor that allows only text input with no GUI support), your shell may be all that you have. You may be used to a windowing environment where you have a lot of programs active at the same time so that you can switch among them as needed. This shell thing can seem pretty limited.

Although the bash shell doesn't include a GUI for running many programs, it does let you move active programs between the background and foreground. In this way, you can have a lot of stuff running, while selectively choosing the one you want to deal with at the moment.
There are several ways to place an active program in the background. One mentioned earlier is to add an ampersand (`&`) to the end of a command line. Another way is to use the `at` command to run commands in a way in which they are not connected to the shell.
To stop a running command and put it in the background, press Ctrl+z. After the command is stopped, you can either bring it to the foreground to run (the `fg` command) or start it running in the background (the `bg` command).

Starting background processes
If you have programs that you want to run while you continue to work in the shell, you can place the programs in the background. To place a program in the background at the time you run the program, type an ampersand (`&`) at the end of the command line. For example:

$ find /usr > /tmp/allusrfiles &
This command finds all files on your Red Hat Linux system (starting from `/usr`), prints those filenames, and puts those names in the file `/tmp/allusrfiles`. The ampersand (`&`) runs that command line in the background. To check which commands you have running in the background, use the `jobs` command, as follows:

$ **jobs**

[1] Stopped (tty output) vi /tmp/myfile

[2] Running find /usr -print > /tmp/allusrfiles &

[3] Running nroff -man /usr/man2/* >/tmp/man2 &

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
[4]- Running        nroff -man /usr/man3/* >/tmp/man3 &
[5]+ Stopped        nroff -man /usr/man4/* >/tmp/man4
```

The first job shows a text-editing command (`vi`) that I placed in the background and stopped by pressing Ctrl+z while I was editing. Job two shows the `find` command I just ran. Jobs three and four show `nroff` commands currently running in the background. Job five had been running in the shell (foreground) until I decided too many processes were running and pressed Ctrl+z to stop job five until a few processes had completed.

The plus sign (+) next to number 5 shows that it was most recently placed in the background. The minus sign (-) next to number 4 shows that it was placed in the background just before the most recent background job. Because job 1 requires terminal input, it cannot run in the background. As a result, it is `Stopped` until it is brought to the foreground again.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip To see the process ID for the background job, add an -l option to the jobs command. If you type ps, you can use the process ID to figure out which command is for a particular ba

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using foreground and background commands

Continuing with the example, you can bring any of the commands on the jobs list to the foreground. For example, to edit `myfile` again, type:

$ **fg %1**

As a result, the `vi` command opens again, with all text as it was when you stopped the vi job.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution Before you put a text processor, word processor, or similar program in the background, make sure you save your file. It's easy to forget you have a program in the background and

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To refer to a background job (to cancel or bring it to the foreground), use a percent sign (%) followed by the job number. You can also use the following to refer to a background job:

% — A percent sign alone refers to the most recent command put into the background (indicated by the plus sign). This action brings the command to the foreground.

%*string* — Refers to a job where the command begins with a particular *string* of characters. The *string* must be unambiguous. (In other words, typing `%vi` when there are two `vi` commands in the background results in an error message.)

%?*string* — Refers to a job where the command line contains a `string` at any point. The string must be unambiguous or the match will fail.

%-- — Refers to the previous job stopped before the one most recently stopped.

If a command is stopped, you can start it running again in the background using the `bg` command. For example, take job number 5 from the jobs list in the previous example:

```
[5]+ Stopped        nroff -man man4/* >/tmp/man4
```

Type the following:

```
$ bg %5
```

After that, the job runs in the background. Its jobs entry appears as follows:

```
[5]  Running        nroff -man man4/* >/tmp/man4 &
```

Configuring your shell

You can tune your shell to help you work more efficiently. Your prompt can provide pertinent information each time you press Enter. You can set aliases to save your keystrokes and permanently set environment variables to suit your needs. To make each change occur when you start a shell, you can add this information to your shell configuration files.

Several configuration files support how your shell behaves. Some of the files are executed for every user and every shell. Others are specific to the user that creates the configuration file. Here are the files that are of interest to anyone using the bash shell in Linux:

`/etc/profile` — This file sets up user environment information for every user. It is executed when you first log in. This file provides values for your path, as well as setting environment variables for such things as the location of your mailbox and the size of your history files. Finally, `/etc/profile` gathers shell settings from configuration files in the `/etc/profile.d` directory.

`/etc/bashrc` — This file is executed for every user that runs the bash shell, each time a bash shell is opened. It sets the default prompt and may add one or more aliases. Values in this file can be overridden by information in each user's `~/.bashrc` file.

`~/.bash_profile` — This file is used by each user to enter information that is specific to their own use of the shell. It is executed only once, when the user logs in. By default it sets a few environment variables and executes the user's `.bashrc` file.

`~/.bashrc` — This file contains the information that is specific to your bash shells. It is read when you log in and also each time you open a new bash shell. This is the best location to add environment variables and aliases so that your shell picks them up.

`~/.bash_logout` — This file executes each time you log out (exit the last bash shell). By default, it simply clears your screen.

To change the `/etc/profile` or `/etc/bashrc` files, you must be the root user. Any user can change the information in the `$HOME/.bash_profile`, `$HOME/.bashrc`, and `$HOME/.bash_logout` files in their own home directories.

The following sections provide ideas about things to add to your shell configuration files. In most cases, you add these values to the `.bashrc` file in your home directory. However, if you administer a system, you may want to set some of these values as defaults for all of your Linux system's users.

Setting your prompt

Your prompt consists of a set of characters that appear each time the shell is ready to accept a command. The `PS1` environment variable sets what the prompt contains. If your shell requires

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

additional input, it uses the values of `PS2`, `PS3`, and `PS4`.

When your Red Hat Linux system is installed, your prompt is set to include the following information: your user name, your host name, and the base name of your current working directory. That information is surrounded by brackets and followed by a dollar sign (for regular users) or a pound sign (for the root user). Here is an example of that prompt:

[chris@myhost bin]$

If you change directories, the bin name would change to the name of the new directory. Likewise, if you were to log in as a different user or to a different host, that information would change.

You can use several special characters (indicated by adding a backslash to a variety of letters) to include different information in your prompt. These can include your terminal number, the date, and the time, as well as other pieces of information. Here are some examples:

`\!` — Shows the current command history number. This includes all previous commands stored for your user name.
`\#` — Shows the command number of the current command. This includes only the commands for the active shell.
`\$` — Shows the user prompt (`$`) or root prompt (`#`), depending on which user you are.
`\W` — Shows only the current working directory base name. For example, if the current working directory was `/var/spool/mail`, this value would simply appear as `mail`.
`\[` — Precedes a sequence of nonprinting characters. This could be used to add a terminal control sequence into the prompt for such things as changing colors, adding blink effects, or making characters bold. (Your terminal determines the exact sequences available.)
`\]` — Follows a sequence of nonprinting characters.
`\\` — Shows a backslash.
`\d` — Displays the day, month, and number of the date. For example: Sat Jan 23.
`\h` — Shows the host name of the computer running the shell.
`\n` — Causes a newline to occur.
`\`*nnn* — Shows the character that relates to the octal number replacing *nnn*.
`\s` — Displays the current shell name. For the bash shell, the value would be `bash`.
`\t` — Prints the current time in hours, minutes, and seconds (for example, 10:14:39).
`\u` — Prints your current user name.
`\w` — Displays the full path to the current working directory.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you are setting your prompt temporarily by typing at the shell, you should put the value of PS1 in quotes. For example, you could type export PS1="[\t \w] \$ " to see a prompt the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To make a change to your prompt permanent, add the value of `PS1` to your `.bashrc` file in your home directory (assuming that you are using the bash shell). There may already be a `PS1` value in that file that you can modify.

Adding environment variables

You may consider adding a few environment variables to your `.bashrc` file. These can help make working with the shell more efficient and effective:

`TMOUT` — This sets how long the shell can be inactive before bash automatically exits. The value is the number of seconds for which the shell has not received input. This can be a nice security feature, in case you leave your desk while you are still logged in to Linux. So as not to be logged you off while you are working, you may want to set the value to something like `TMOUT=1800` (to allow 30 minutes of idle time).

`PATH` — As described earlier, the `PATH` variable sets the directories that are searched for commands you use. If you often use directories of commands that are not in your `PATH`, you can permanently add them. To do this, add a `PATH` variable to your `.bashrc` file. For example, to add a directory called `/getstuff/bin`, add the following:

```
PATH=$PATH:/getstuff/bin ; export PATH
```

This example first reads all the current path directories into the new `PATH` (`$PATH`), adds the `/getstuff/bin` directory, and then exports the new `PATH`.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

Some people add the current directory to their PATH by adding a directory identified simply as a dot (.), as follows:

```
PATH=.:$PATH;export PATH
```

This lets you always run commands in your current directory (which people may be used to if they have used DOS). However, Howe

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

WHATEVER — You can create your own environment variables to provide shortcuts in your work. Choose any name that is not being used and assign a useful value to it. For example, if you do a lot of work with files in the `/work/time/files/info/memos` directory, you could set the following variable:

```
M=/work/time/files/info/memos ; export M
```

You could make that your current directory by typing `cd $M`. You could run a program from that directory called hotdog by typing `$M/hotdog`. You could edit a file from there called bun by typing `vi $M/bun`.

Adding aliases

Setting aliases can save you even more typing than setting environment variables. With aliases, you can have a string of characters execute an entire command line. You can add and list aliases with the `alias` command. Here are some examples:

```
alias p='pwd ; ls -CF'

alias rm='rm -i'
```

In the first example, the letter `p` is assigned to run the command `pwd`, and then to run `ls -CF` to print the current working directory and list its contents in column form. The second runs the `rm` command with the `-i` option each time you simply type `rm`. (This is an alias that is often set automatically for the root user, so that instead of just removing files, you are prompted for each individual file removal. This prevents you from removing all the files in a directory by mistakenly typing something such as `rm *`.) While you are in the shell, you can check which aliases are set by typing the `alias` command. If you want to remove an alias, you can type `unalias`. (Remember that if the `alias` is set in a configuration file, it will be set again when you open another shell.)

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Working with the Red Hat Linux File System

The Red Hat Linux file system is the structure in which all the information on your computer is stored. Files are organized within a hierarchy of directories. Each directory can contain files, as well as other directories.

If you were to map out the files and directories in Red Hat Linux, it would look like an upside-down tree. At the top is the root directory, which is represented by a single slash (`/`). Below that is a set of common directories in the Linux system, such as `/bin`, `/dev`, `/home`, `/lib`, and `/tmp`, to name a few. Each of those directories, as well as directories added to the root, can contain subdirectories. Figure 4-1 illustrates how the Linux file system is organized as a hierarchy. To illustrate how directories are connected, Figure 4-1 shows a `/home` directory that contains subdirectories for three users: `chris`, `mary`, and `tom`. Within the `chris` directory are subdirectories: `briefs`, `memos`, and `personal`. To refer to a file called `inventory` in the `chris/memos` directory, you could type the full path of `/home/chris/memos/inventory`. If your current directory were `/home/chris/memos`, you could refer to the file as simply `inventory`.

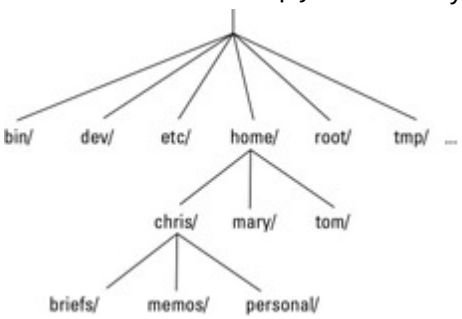

Figure 4-1: The Red Hat Linux file system is organized as a hierarchy of directories.

Some of the Red Hat Linux directories that may interest you include the following:

- `/bin` — Contains common Linux user commands, such as `ls`, `sort`, `date`, and `chmod`.
- `/boot` — Has the bootable Linux kernel and boot loader configuration files (GRUB).
- `/dev` — Contains files representing access points to devices on your systems. These include terminal devices (`tty*`), floppy disks (`fd*`), hard disks (`hd*`), RAM (`ram*`), and CD-ROM (`cd*`). (Users normally access these devices directly through the device files.)
- `/etc` — Contains administrative configuration files.
- `/home` — Contains directories assigned to each user with a login account.
- `/mnt` — Provides a location for mounting devices, such as remote file systems and removable media (with directory names of cdrom, floppy, and so on).
- `/root` — Represents the root user's home directory.
- `/sbin` — Contains administrative commands and daemon processes.
- `/tmp` — Contains temporary files used by applications.
- `/usr` — Contains user documentation, games, graphical files (X11), libraries (lib), and a variety of other user and administrative commands and files.
- `/var` — Contains directories of data used by various applications. In particular, this is where you would place files that you share as an FTP server (`/var/ftp`) or a Web server (`/var/www`). It also contains all system log files (`/var/log`).

The file systems in the DOS or MS Windows operating systems differ from Linux's file structure. See the sidebar on the Linux file system versus MS file systems.

Creating files and directories

As a Red Hat Linux user, most of the files you save and work with will probably be in your home directory. Here are commands you use to create and use files and directories:

- `cd` — Change to another current working directory
- `pwd` — Print the name of the current working directory
- `mkdir` — Create a directory

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

`chmod` — Change the permission on a file or directory
`ls` — List the contents of a directory

The following procedure steps you through creating directories within your home directory, moving among your directories, and setting appropriate file permissions:

1. First, go to your home directory. To do this, simply type `cd`. (For other ways of referring to your home directory, see the sidebar on identifying directories.)

Red Hat Linux File System versus Windows-based File Systems

3. [redacted]milar in many ways, the Linux file system has some striking differences [redacted]tems used in MS-DOS and Windows operating systems. Here are a

   [redacted]MS-DOS and Windows file systems, drive letters represent different storage devices (for example, A: is a floppy drive and C: is a hard disk). In Linux, all storage devices are fit into the file system hierarchy. So, the fact that all of `/usr` may be on a separate hard disk or that `/mnt/rem1` is a file system from another computer is invisible to the user.
   Slashes, rather than backslashes, are used to separate directory names in Linux. So, `C:\home\chris` in an MS system is `/home/chris` in a Linux system.
   Filenames almost always have suffixes in DOS (such as .txt for text files or .doc for word-processing files). Although at times you can use that convention in Linux, three-character suffixes have no required meaning in Linux. They can be useful for identifying a file type.
   Every file and directory in a Linux system has permissions and ownership associated with it. Security varies among Microsoft systems. Because DOS and MS Windows began as single-user systems, file ownership was not built into those systems when they were designed. Later releases added features such as file and folder attributes to address this problem.

4. To make sure that you got to your home directory, type `pwd`. When I do this, I get the following response (yours will reflect your home directory):
5. $ **pwd**
   /home/chris
6. Create a new directory called `test` in your home directory, as follows:
7. $ **mkdir test**
8. Check the permissions of the directory by typing:
9. $ **ls -ld test**
   drwxr-xr-x 2 chris sales 1024 Jan 24 12:17 test
   Notice that this listing says that test is a directory (`d`), the owner is `chris`, the group is `sales`, and the file was most recently modified on Jan 24 at 12:17 p.m. Suppose that you want to prevent everyone else who uses this computer from using or viewing the files in this directory. The permissions for the directory are `rwxr-xr-x`. I explain what these permissions mean later in this section.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note When you add a new user in Red Hat Linux, by default, the user is assigned to a group of the same name. For example, in the preceding text, the user chris would be assigned to

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

10. For now, type the following:
11. $ **chmod 700 test**
    This step changes the permissions of the directory to give you complete access and everyone else no access at all. (The new permissions should read like `rwx------`.)
12. Next, make the test directory your current directory as follows:
13. $ **cd test**

Identifying Directories

When you need to identify your home directory on a shell command line, you can use the following:

`$HOME` — This environment variable stores your home directory name.
~ — The tilde (~) represents your home directory on the command line.

You can also use the tilde to identify someone else's home directory. For example, `~chris` would be expanded to the chris home directory (probably `/home/chris`).

Other special ways of identifying directories in the shell include the following:

`.` — A single dot (`.`) refers to the current directory.
`..` — Two dots (`..`) refers to a directory directly above the current directory.
`$PWD` — This environment variable refers to the current working directory.
`$OLDPWD` — This environment variable refers to the previous working directory before you changed to the current one.

Using metacharacters and operators

To make more efficient use of your shell, the bash shell lets you use certain special characters, referred to as metacharacters and operators. Metacharacters can help you match one or more files without typing each file completely. Operators let you direct information from one command or file to another command or file.

Using file-matching metacharacters

To save you some keystrokes and to be able to refer easily to a group of files, the bash shell lets you use metacharacters. Anytime you need to refer to a file or directory, such as to list it, open it, or remove it, you can use metacharacters to match the files you want. Here are some useful metacharacters for matching filenames:

***** — This matches any number of characters.
? — This matches any one character.
`[...]` — This matches any one of the characters between the brackets, which can include a dash-separated range of letters or numbers.

To try out some of these file-matching metacharacters, go to an empty directory (such as the `test` directory described in the previous section) and create some files. Here's an example of how to create some empty files:

$ **touch apple banana grape grapefruit watermelon**

The next few commands show you how to use shell metacharacters to match filenames so they can be used as arguments to the `ls` command. Using metacharacters shown below, you can match the filenames you just created with the `touch` command. Type the following commands and see if you get the same responses:

$ **ls a***

apple

$ **ls g***

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

grape

grapefruit

$ **ls g*t**

grapefruit

$ **ls *e***

apple grape grapefruit watermelon

$ **ls *n***

banana watermelon

The first example matches any file that begins with an `a` (`apple`). The next example matches any files that begin with `g` (`grape`, `grapefruit`). Next, files beginning with `g` and ending in `t` are matched (`grapefruit`). Next, any file that contains an `e` in the name is matched (`apple`, `grape`, `grapefruit`, `watermelon`). Finally, any file that contains an `n` is matched (`banana`, `watermelon`).

Here are a few examples of pattern matching with the question mark (`?`):

$ **ls ????e**

apple grape

$ **ls g???e***

grape grapefruit

The first example matches any five-character file that ends in `e` (`apple`, `grape`). The second matches any file that begins with `g` and has `e` as its fifth character (`grape`, `grapefruit`).

Here are a few examples of using braces to do pattern matching:

$ **ls [abw]***

apple banana watermelon

$ **ls [agw]*[ne]**

apple grape watermelon

In the first example, any file beginning with `a`, `b`, or `w` is matched. In the second, any file that begins with `a`, `g`, or `w` and also ends with either `n` or `e` is matched.

Using file-redirection metacharacters

Commands receive data from standard input and send it to standard output. Using pipes (described earlier), you can direct standard output from one command to the standard input of another. With files, you can use less than (<) and greater than (>) signs to direct data to and from files. Here are the file redirection characters:

- **<** — Direct the contents of a file to the command.
- **>** — Direct the output of a command to a file, deleting the existing file.
- **>>** — Direct the output of a command to a command, adding the output to the end of the existing file.

Here are some examples of command lines where information is directed to and from files:

$ **mail root < ~/.bashrc**

$ **man chmod | col -b > /tmp/chmod**

$ echo "I finished the project on $(date)" >> ~/projects

In the first example, the contents of the `.bashrc` file in the home directory are sent in a mail message to the computer's root user. The second command line formats the chmod man page (using the `man` command), removes extra back spaces (`col -b`) and sends the output to the file `/tmp/chmod` (erasing the previous `/tmp/chmod` file, if it exists). The final command results in the following text being added to the user's project file:

I finished the project on Sat Oct 25 13:46:49 PST 2003

Understanding file permissions

After you've worked with Linux for a while, you are almost sure to get a "Permission denied" message.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Permissions associated with files and directories in Linux were designed to keep users from accessing other users' private files and to protect important system files.
The nine bits assigned to each file for permissions define the access that you and others have to your file. Permission bits appear as `rwxrwxrwx`. The first three bits apply to the owner's permission, the next three apply to the group assigned to the file, and the last three apply to all others. The `r` stands for read, the `w` stands for write, and the `x` stands for execute permissions. If a dash appears instead of the letter, it means that permission is turned off for that associated read, write, or execute.
You can see the permission for any file or directory by typing the `ls -ld` command. The named file or directory appears as those shown in the following example:

$ **ls -ld ch3 test**

-rw-rw-r-- 1 chris sales 4983 Jan 18 22:13 ch3

drwxr-xr-x 2 chris sales 1024 Jan 24 13:47 test
The first line shows a file (`ch3`) that has read and write permission on for the owner and the group. All other users have read permission, which means they can view the file but cannot change its contents or remove it. The second line shows a directory (indicated by the letter d before the permission bits). The owner has read, write, and execute permission, while the group and other users have only read and execute permissions. As a result, only the owner can add, change, or delete files in that directory. Any other user, however, can only read the contents, change to that directory, and list the contents of the directory.
If you own a file, you can change the permission on it as you please. You can do this with the `chmod` command. For each of the three sets of permission on a file (read, write, and execute), the r is assigned to the number 4, w to 2, and x to 1. So to make permissions wide open for yourself as owner, you would set the first number to 7 (4 plus 2 plus 1). The same would be true for group and other permission. Any combination of permissions can result from 0 (no permission) through 7 (full permission).

Here are some examples of how to change permission on a file and what the resulting permission would be:

chmod 777 files (rwxrwxrwx

chmod 755 files (rwxr-xr-x

chmod 644 files (rw--r—r-

chmod 000 files (---------
When you try to create a file, by default it is given the permission rw-r--r--. A directory is given the permission rwxr-xr-x. These default values are determined by the value of `umask`. Type `umask` to see what your `umask` value is. For example:

$ **umask**

022

The umask value masks the permissions value of 666 for a file and 777 for a directory. The umask of 022 results in permission for a directory of 755 (rwxr-xr-x). That same umask results in a file permission of 644 (rw-r--r--). (Execute permissions are off by default for regular files.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Here's a great tip for changing the permission for lots of files at once. Using the -R options of chmod, you could change the permission for all of the files and directories within a directory structure at once. For example, if you wanted to open permissions completely to all files and directories in the /tmp/test directory

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution The -R option of chmod works best if you are opening permissions completely or adding execute permission (as well as the appropriate read/write permission). The reason is that if

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Moving, copying, and deleting files

Commands for moving, copying, and deleting files are fairly straightforward. To change the location of a file, use the `mv` command. To copy a file from one location to another, use the `cp` command. To remove a file, use the `rm` command. Here are some examples:

$ **mv abc def**

$ **mv abc ~**

$ **cp abc def**

$ **cp abc ~**

$ **rm abc**

$ rm *

Of the two move (`mv`) commands, the first moves the file `abc` to the file `def` in the same directory (essentially renaming it), whereas the second moves the file `abc` to your home directory (~). The first copy command (`cp`) copies `abc` to the file `def`, whereas the second copies `abc` to your home directory (~). The first remove command (`rm`) deletes the `abc` file; the second removes all the files in the current directory (except those that start with a dot).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note For the root user, the mv, cp, and rm commands are aliased to each be run with the -i option. This causes a prompt to appear asking you to confirm each move, copy, and remo

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the vi Text Editor

It's almost impossible to use Red Hat Linux for any period of time and not need to use a text editor. If you are using a GUI, you can run `gedit`, which is fairly intuitive for editing text. Most Red Hat Linux shell users will use either the `vi` or `emacs` commands to edit text files. The advantage of `vi` or `emacs` over a graphical editor is that you can use it from any shell, a character terminal, or a character-based connection over a network (using `telnet` or `ssh`, for example). No GUI is required.

This section provides a brief tutorial of the `vi` text editor. Any time in this book that I suggest you manually edit a configuration file, you can use `vi` to do that editing (from any shell). (If `vi` doesn't suit you, see the "Exploring Other Text Editors" sidebar for other options.)

The `vi` editor is difficult to learn at first. But when you know it, you will be able to edit and move around quickly and efficiently within files. Your fingers never have to leave the keyboard to pick up a mouse or hit a function key.

[redacted]

Exploring Other Text Editors

There [redacted] of text editors available to use with Linux. Here are a few contained in Red Hat Linux that you can try out if you find vi to be too taxing:

gedit — The GNOME text editor that runs in the GUI.
jed — This screen-oriented editor was made for programmers. Using colors, jed can highlight code you create so you can easily read the code and spot syntax errors. Use the Alt key to select menus to manipulate your text.
joe — The joe editor is similar to many PC text editors. Use control and arrow keys to move around. Type Ctrl+C to exit with no save or Ctrl+X to save and exit.
kedit — A GUI-based text editor that comes with the KDE desktop.
mcedit — With mcedit, function keys help you get around, save, copy, move, and delete text. Like jed and joe, mcedit is screen-oriented.

If you use ssh to log in to other Linux computers on your network, you can use any editor to edit files. A GUI-based editor will pop up on your screen. When no GUI is available, you will need a text editor that runs in the shell, such as vi, jed or joe.

[redacted]

Starting wi

Most often, [redacted] to open a particular file. For example, to open a file called `/tmp/test`, type the following command:

$ **vi /tmp/test**

If this is a new file, you should see something similar to the following:

~
~
~
~
~
"/tmp/test" [New File]

The box at the top represents where your cursor is. The bottom line keeps you informed about what is going on with your editing (here you just opened a new file). In between, there are tildes (~) as filler because there is no text in the file yet. Now here's the intimidating part: there are no hints, menus, or icons to tell you what to do. On top of that, you can't just start typing. If you do, the computer is likely to beep at you.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The first things you need to know are the different operating modes. The `vi` editor operates in either command mode or input mode. When you start `vi`, you are in command mode. Before you can add or change text in the file, you have to type a command to tell `vi` what you want to do. A command consists of one or two letters and an optional number. To get into input mode, you need to type an input command. To start out, type either of the following input commands:

a — Add. After you type `a`, you can input text that starts to the right of the cursor.
i — Insert. After you type `i`, you can input text that starts to the left of the cursor.

Type a few words and press Enter. Repeat that a few times until you have a few lines of text. When you are done typing, press Esc. You are now back in command mode. Now that you have a file with some text in it, try moving around within that text with the following keys or letters:

Tip Remember the Esc key! It always places you back into command mode.

Arrow keys — Use the arrow keys to move up, down, left, or right in the file one character at a time. To move left and right you can also use Backspace and the Spacebar, respectively. If you prefer to keep your fingers on the keyboard, use h (left), l (right), j (down), or k (up) to move the cursor.
w — Moves the cursor to the beginning of the next word.
b — Moves the cursor to the beginning of the previous word.
0 (*zero*) — Moves the cursor to the beginning of the current line.
$ — Moves the cursor to the end of the current line.
H — Moves the cursor to the upper-left corner of the screen (first line on the screen).
M — Moves the cursor to the first character of the middle line on the screen.
L — Moves the cursor to the lower-left corner of the screen (last line on the screen).

Now that you know how to input text and move around, the only other editing you need to know is how to delete text. Here are a few `vi` commands for deleting text:

x — Deletes the character under the cursor.
X — Deletes the character directly before the cursor.
dw — Deletes from the current character to the end of the current word.
d$ — Deletes from the current character to the end of the current line.
d0 — Deletes from the previous character to the beginning of the current line.

If you feel pretty good about creating text and moving around the file, you may want to wrap things up. Several keystrokes for saving and quitting the file follow:

ZZ — Save the current changes to the file and exit from vi.
:w — Save the current file but continue editing.
:wq — Same as ZZ.
:q — Quit the current file. This works only if you don't have any unsaved changes.
:q! — Quit the current file and *don't* save the changes you just made to the file.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you've really trashed the file by mistake, the :q! command is the best way to exit and abandon your changes. The file reverts to the most recently changed version. So, if you

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You have learned a few `vi` editing commands. I describe more commands in the following sections. However, before I do, here are a few tips to smooth out your first trials with `vi`:

Esc — Remember that Esc gets you back to command mode. (I've watched people press every key on the keyboard trying to get out of a file.) Esc followed by ZZ gets you out of command mode, saves the file, and exits.

u — Press `u` to undo the previous change you made. Continue to press `u` to undo the change before that, and the one before that.

Ctrl+r — If you decide you didn't want to undo the previous command, use Ctrl+r for Redo. Essentially, this command undoes your undo.

Caps Lock — Beware of hitting Caps Lock by mistake. Everything you type in `vi` has a different meaning when the letters are capitalized. You don't get a warning that you are typing capitals — things just start acting weird.

:! *command* — You can run a command while you are in `vi` using `:!` followed by a command name. For example, type **:!date** to see the current date and time, type **:!pwd** to see what your current directory is, or type **:!jobs** to see if you have any jobs running in the background. When the command completes, press Enter and you are back to editing the file. You could even do that with a shell (`:!bash`) to run a few commands from the shell, then type **exit** to return to `vi`. (I recommend doing a save before escaping to the shell, just in case you forget to go back to `vi`.)

-- INSERT -- — When you are in insert mode, the word INSERT appears at the bottom of the screen. Other messages also appear at the line at the bottom of the screen.

Ctrl+g — If you forget what you are editing, pressing these keys displays the name of the file that you are editing and the current line that you are on. It also displays the total number of lines in the file, the percentage of how far you are through the file, and the column number the cursor is on. This just helps you get your bearings after you've stopped for a cup of coffee at 3a.m.

Moving around the file

Besides the few movement commands described earlier, there are other ways of moving around a `vi` file. To try these out, open a large file that you can't do much damage to. (Try copying `/var/log/messages` to `/tmp` and opening it in vi.) Here are some movement commands you can use:

Ctrl+f — Page ahead, one page at a time.

Ctrl+b — Page back, one page at a time.

Ctrl+d — Page ahead 1/2 page at a time.

Ctrl+u — Page back 1/2 page at a time.

G — Go to the last line of the file.

1G — Go to the first line of the file. (Use any number to go to that line in the file.)

Searching for text

To search for the next occurrence of text in the file, use either the slash (`/`) or the question mark (`?`) character. Within the search, you can also use metacharacters. Here are some examples:

`/hello` — Searches forward for the word `hello`.

`?goodbye` — Searches backward for the word `goodbye`.

`/The.*foot` — Searches forward for a line that has the word `The` in it and also, after that at some point, the word `foot`.

`?[pP]rint` — Searches backward for either `print` or `Print`. Remember that case matters in Linux, so using brackets can search for words that could have different capitalization.

The `vi` editor was originally based on the `ex` editor. That editor did not let you work in full-screen mode. However, it did enable you to run commands that let you find and change text on one or more lines at a time. When you type a colon and the cursor goes to the bottom of the screen, you are essentially in `ex` mode. Here is an example of some of those `ex` commands for searching for and changing text. (I chose the words `Local` and `Remote` to search for, but you can use any appropriate words.)

`:g/Local` — Searches for the word `Local` and prints every occurrence of that line from the file. (If there is more than a screenful, the output is piped to the `more` command.)

`:s/Local/Remote` — Substitutes `Remote` for the word `Local` on the current line.

`:g/Local/s//Remote` — Substitutes the first occurrence of the word `Local` on

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

every line of the file with the word `Remote`.

`:g/Local/s//Remote/g` — Substitutes every occurrence of the word `Local` with the word `Remote` in the entire file.

`:g/Local/s//Remote/gp` — Substitutes every occurrence of the word `Local` with the word `Remote` in the entire file, then prints each line so that you can see the changes (piping it through `more` if output fills more than one page).

Using numbers with commands

You can precede most `vi` commands with numbers to have the command repeated that number of times. This is a handy way to deal with several lines, words, or characters at a time. Here are some examples:

`3dw` — Deletes the next three words.

`5cl` — Changes the next five letters (that is, removes the letters and enters input mode).

`12j` — Moves down 12 lines.

Putting a number in front of most commands just repeats those commands. At this point, you should be fairly proficient at using the `vi` command.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

Working from a shell command-line interpreter within Red Hat Linux may not be as simple as using a GUI, but it offers many powerful and flexible features. This chapter describes how to log in to Red Hat Linux and use shell commands. Features for running commands include recalling commands from a history list, completing commands, and joining commands.

This chapter describes how shell environment variables can be used to store and recall important pieces of information. It also teaches you to modify shell configuration files to tailor the shell to suit your needs. Finally, the chapter describes how to use the Red Hat Linux file system to create files and directories, use permissions, and work with files (moving, copying, and removing them), and how to edit text files from the shell using the `vi` command.

Team LiB

PREVIOUS NEXT

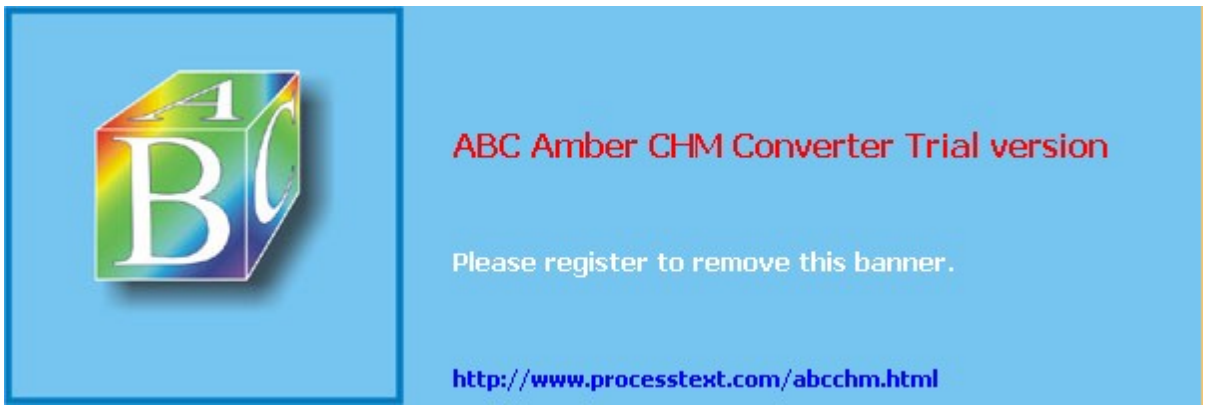

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 5: Accessing and Running Applications

Overview

To get your work done on a computer, you use application programs. They let you create documents, crunch data, and communicate with others. As an engine for running applications, Red Hat Linux is becoming more viable every day. Not too long ago, there were only a handful of user-friendly applications available. Now there are hundreds — and they're getting more powerful and friendlier all the time.

This chapter describes how to get applications for Red Hat Linux and run them. In particular, it teaches you how to download applications from the Internet and to uncompress and install them. For running native Linux applications, this chapter focuses on graphical-based applications (run on an X desktop, including applications distributed over the network).

Besides programs that were specifically created for (or ported to) Linux, it is possible to run applications that were intended for other operating systems. This chapter describes emulators that are available to use with Red Hat Linux for running applications created for Windows, DOS, and Macintosh operating systems.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using Red Hat Linux as an Application Platform

Although operating systems are nice (and necessary), people use desktop computers to run application programs. A strong case can be made for using Red Hat Linux as a server, but as a desktop system, Red Hat Linux is still some distance away from challenging Microsoft's operating system dominance for several reasons:

- Although you can get word processing programs, spreadsheet programs, graphics programs, and almost any other type of application that you want for Linux, many of the most popular applications in each category don't run well in Linux or don't run at all. For example, the latest Microsoft Office product will not run in Linux (though `Codeweavers.com` and others have gotten earlier versions of Microsoft Office to run in Linux). If your company uses MS Word for word processing or MS Excel for spreadsheets, you could try using those files with StarOffice or OpenOffice in Red Hat Linux. However, those files won't always convert cleanly.
- There are many more commercial, battle-tested desktop applications for Microsoft Windows operating systems than there are for Linux. Because the market is so huge for Windows systems, many software companies develop their products solely for that market.
- Linux applications, as a rule, have historically been more difficult to configure and use than many commercial Windows applications.

That's the bad news. The good news is that Linux is significantly gaining ground. You can now use Linux on your desktop to do almost everything you would want to do on a desktop computer with Windows. The time is coming when you will be able to replace the operating systems on the desktop computers in your home or office with Red Hat Linux. I believe that, in the long term, Linux could become a preferred operating system for running applications. Here are some reasons:

- Many people believe that networked applications will drive the future of computing. Unlike the first Windows systems, which had their roots in the single-user, one-task-at-a-time DOS system, Red Hat Linux is based on UNIX systems. UNIX was designed from the ground up to deal with many users and many tasks in a networked environment. Red Hat Linux offers a strong foundation for networked applications.
- A huge development community is working on open source applications to meet the needs of the Linux community. Recently, some strong commercial offerings have been added.
- In the spirit of Linux and the GNU (which stands for "*GNU is Not UNIX*"), most application programs are free or inexpensive. This means that you can try out most applications for little or no money. Getting started running Linux applications can be done at a small cost. For example, OpenOffice (the open source version of StarOffice) comes with Red Hat Linux and the boxed set of StarOffice is only $75.95 (from `www.sun.com/staroffice`), whereas Microsoft Office will cost you hundreds of dollars. Although you can buy components such as Microsoft Word separately, even they can cost several hundred dollars apiece.
- The transition for desktop users from Windows to Linux is becoming easier, with great strides being made by projects such as the WINE project. WINE (described later in this chapter) lets you run many Windows programs directly in Linux. WINE also provides a path for application developers to more simply port their Windows code to run in Linux.

The bottom line is that it will take some effort for most people to discard their Windows operating systems completely. However, if you are committed to making Red Hat Linux your sole application platform, there are several ways to ease that transition. Emulation programs let you run many programs that were created for other operating systems. Conversion programs can help you convert graphics and word processing data files from other formats to those supported by Linux applications.

Cross-Reference See Chapter 6 for information on importing and exporting word processing and graphics files.

If you are running Linux on a PC, chances are that you already paid for a Windows 95, 98, ME, XP, NT, or 2000 operating system. You can either run the different operating systems on different PCs or have Windows and Linux on separate partitions of your hard disk. The latter requires that you reboot each

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

time you want to switch operating systems.

The following section describes applications that run in Red Hat Linux that you can use to replace the Windows applications you are used to.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Finding Windows-Equivalent Applications in Linux

If you are going to use Red Hat Linux as a desktop computer system, you have to be able to write documents, work with graphics, and crunch numbers. You probably also have other favorite applications, like a music player, Web browser, and e-mail reader.

Note Using WINE technology, the people at Codeweavers, Inc. offer a CrossOver Office product that lets you install and run Microsoft Office in Linux. See the "Running Windows Applications with WINE" section later in this chapter.

To give you a snapshot of what desktop applications are available, Table 5-1 contains a list of popular Windows applications, equivalent Linux applications, and where you can find the Linux applications. Although many of these applications have not reached the level of sophistication of their Windows counterparts, they can be cost-effective alternatives.

Table 5-1: Windows-Equivalent Linux Applications

Windows Applications	Linux Applications	Where to Get Linux Applications	Cost
Microsoft Office (office productivity suite)	OpenOffice (`openoffice.org`) Koffice StarOffice	Included on Red Hat CDs Included on Red Hat CDs `www.sun.com/staroffice`	Free Free $75.95
Microsoft Word (word processor)	OpenOffice Writer abiWord kword	Included on Red Hat CDs Included on Red Hat CDs Included on Red Hat CDs	Free Free Free
Microsoft Excel (spreadsheet)	OpenOffice Calc gnumeric kspread	Included on Red Hat CDs Included on Red Hat CDs Included on Red Hat CDs	Free Free Free
Microsoft Powerpoint (presentation)	OpenOffice Impress kpresenter	Included On Red Hat CDs Included on Red Hat CDs	Free Free
Microsoft Internet Explorer (Web browser)	mozilla epiphany konqueror opera	Included on Red Hat Linux CDs Included on Red Hat Linux CDs Included on Red Hat Linux CDs `www.opera.com`	Free Free Free Free download

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 5-1: Windows-Equivalent Linux Applications

Windows Applications	Linux Applications	Where to Get Linux Applications	Cost
Microsoft Outlook (e-mail reader)	evolution kmail Mozilla Mail	Included on Red Hat CDs Included on Red Hat CDs Included on Red Hat CDs	Free Free Free
Adobe Photoshop (image editor)	The Gimp (gimp)	Included on Red Hat CDs	Free
Quicken or Microsoft Money	gnucash	Included on Red Hat CDs	Free
AutoCAD (computer-aided design)	LinuxCad	www.linuxcad.com	$89

The following sections describe how to find and work with application programs that are created specifically for Linux.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Obtaining Red Hat Linux Applications

Unfortunately, you won't be able to walk into the average computer store and find a lot of Linux application programs. The best way to get Linux applications (other than those included with your Red Hat Linux system) is to download them from the Internet. They can also be ordered on CD-ROM from several Linux Web sites.

Software packages that are specifically compiled and packaged for Red Hat Linux distributions are almost always available in RPM format. So when you begin scouring the Internet for Red Hat software (as described later), look for software repositories of RPM packages built specifically for the version of Red Hat Linux or Fedora Core that you are using.

Investigating your desktop

More and more high-quality desktop applications are being packaged with Red Hat Linux, mostly as part of the GNOME or KDE desktop environments. In other words, to start finding some excellent office applications, games, multimedia players, and communications tools, you don't have to look any further than the red hat menu button on your desktop.
So before you start hunting around the Internet for the software you need, see if you can use something already installed with Red Hat Linux. The chapters that follow this one describe how to use publishing tools, play games, work with multimedia, and communicate over the Internet — all with programs that are either on the CDs that come with this book or are easily attainable. Appendix B contains a list of the software packages that are included with the complete version of Red Hat Linux that is packaged with this book.

Tip To keep up with fixes to software packages that are part of the Red Hat Linux distribution, Red Hat offers a service for automatically downloading and installing packages. Chapter 10 shows you how to use Red Hat Network and the up2date command to update your Red Hat Linux software.

Finding applications on the Internet

If you don't already know the names of applications you want to use, there are a lot of places to look for Linux applications on the Internet. If you do know what you want, your best bet might be to head for a software repository that has packages build specifically for Red Hat Linux and Fedora.

Here are a few Web sites that you can browse for software that runs in Linux:

Freshmeat (`www.freshmeat.net`) — This site maintains a massive index of Linux software. You can do keyword searches for software projects or browse for software by category.

SourceForge (`www.sourceforge.net`) — This site hosts thousands of open source software projects. You can download software and documentation from those projects from the SourceForge site.

Tucows Linux site (`linux.tucows.com`) — Both free and commercial software for Linux is available from the Tucows Linux Web site. This site also features news articles on Linux and a listing of software downloads from the site by category.

When you purchase a commercial application, you usually get the application on CD. Installation is often simplified, and hard copy documentation is provided. Of course, when you download software, you get immediate gratification — when you purchase it, you have to wait for the package to arrive in the mail.

Tip Sometimes software packages will be available in both libc5 and libc6 formats. This designation refers to the version of C programming-language libraries used by the application. If you have a choice, choose the libc6 packages. These are compatible with Red Hat Linux 7 and later (or any Linux kernel version 2.2 and higher). In fact, all major Linux distributions now use libc6. Better yet, look for packages designated for the specific distribution you are using (such as Fedora, Red Hat Linux 9, and so on.).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You can visit FTP sites containing RPM packages if you already have some idea of what you are looking for. You can start by reading the README and INDEX files for a particular software product to get your bearings. When you find the software you want, simply click the link to begin downloading the software to your computer. Here are a few sites that are particularly good for finding RPM packages:

Rpmfind (`www.rpmfind.net`) — Open source software that is already packaged up in Red Hat Package Management (RPM) format is available from this site. Do a keyword search from this Web site, then download the appropriate RPM from the search results.

FreshRPMs (`www.freshrpms.net`) — Another site with a good selection of high-quality RPMs.

Red Hat FTP Mirrors (`www.redhat.com/mirrors.html`) — Go to this page for a listing of FTP sites containing Red Hat software that you can download. Most of these sites also have a variety of freeware and shareware applications that are usable with Linux.

Often, you can't just download a single software package to get the software in that package to work. Many packages depend on other packages. For example, software packages for playing audio and video typically rely on other software packages for decoding different kinds of content. To deal with this issue, Red Hat Linux (Fedora Core) has included the yum package.

Downloading and installing applications with yum

The Yellow dog Updater, Modified (yum) software package lets you install and update selected software packages in RPM format from software repositories on the Web. Once you know the software package that you want, yum is probably the best way to download and install that package.

The yum package is included on the Fedora Core CDs that come with this book. To use yum to install RPM software packages, follow these basic steps:

1. **Determine the software package you want.** Many popular add-on packages for Red Hat Linux are already built for specific versions Red Hat Linux (Red Hat 8, 9, Fedora, and so on) and stored in software repositories on the Internet. Lists of repositories are available from the Yum Web site at `http://linux.duke.edu/projects/yum/repos`.
2. **Configure yum.** You need to configure the `/etc/yum.conf` file to point to the repository that contains the software you want. Once that is done, you can install any package that repository contains.
3. **Run yum.** The `yum` command can be used to download and install any package from the yum repository, including any packages the one you want depends on.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

In Red Hat Linux, and presumably in the new Fedora Core, Red Hat, Inc. has gone to great lengths to ensure that software it provides is of good quality and unimpaired by legi

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Besides downloading and installing new software packages, yum can also be used to check for available updates and list various kinds of information about available packages.

Configuring yum (/etc/yum.conf)

The `/etc/yum.conf` file already comes preconfigured to be able to install Red Hat base system packages as well as updates. So, to update your system with packages that are part of the current Red Hat Linux release, you probably don't need to update `yum.conf` at all.
To be able to download from a repository that contains software that is not in a Red Hat Linux or Fedora distribution, you need to add that repository to the `/etc/yum.conf` file. Here is an example of an entry in yum.conf that points to such a repository:

[freshrpms-rhl-9]

name=Freshrpms packages for Red Hat Linux 9

baseurl=http://ayo.freshrpms.net/redhat/9/i386/freshrpms
In this example, for the entry `freshrpms-rhl-9`, the repository is named `Freshrpms packages for Red Hat Linux 9`. When you request a software package with yum, it looks in subdirectories of `http://ayo.freshrpms.net/redhat/9/i386/freshrpms`. If you are using Fedora Core, you can check `freshrpms.net` or similar sites for a repository that is specific to Fedora Core.

Running yum to download and install RPMs

With the repository identified in your yum.conf file, downloading and installing an RPM you want is as simple as running `yum` with the `install` option to request the RPM. With an active connection to the Internet, open a Terminal window as root user. Here is an example of using the `yum` command to download the `mplayer` media player:

yum install mplayer

Gathering header information file(s) from server(s)

Server: Red Hat Linux 0.94 - i386 - Base

Server: Freshrpms packages for Red Hat Linux 9

Server: Red Hat Linux 0.94 - Updates

Finding updated packages

Downloading needed headers

Resolving dependencies

..Dependencies resolved

I will do the following:

[install: mplayer 1.0-0.1.20031002.fr.i386]

I will install/upgrade these to satisfy the dependencies:

.

.

.

Is this ok [y/N]: **y**

.

.

.

Getting mplayer-fonts-1.1-1.fr.noarch.rpm

Getting mplayer-1.0-0.1.20031002.fr.i386.rpm

Calculating available disk space - this could take a bit

mplayer-fonts 100 % done 2/14

mplayer 100 % done 14/14

Installed: mplayer 1.0-0.1.20031002.fr.i386

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Dep Installed: mplayer-fonts 1.1-1.fr.noarch

Transaction(s) Complete
As you can see from this example, `yum` checked three different software repositories: Red Hat Linux 0.94 (Base and Updates) and Freshrpms packages for Red Hat Linux 9. After listing the dependencies, yum asks if it is OK to install them. Type **`y`** and the package and all its dependencies are installed.

Using yum for listing and updating packages

Besides downloading and installing new RPM packages, yum can also be used to list available packages and update packages that are already installed. The following examples illustrate some uses of yum.

yum check-update
The `check-update` option causes yum to check the software repositories for available updated versions of RPM packages you have installed. If you see a package you want to update, you can use the `update` option. For example, to update the nmap-frontend package, you could type the following:

yum update nmap-frontend

If you want to see a list of all packages that are available for download from the repositories you have entered, type the following:

yum list | less
Adding the `less` command to the end lets you page through the list of software (it could be long, depending on which repositories you point to). If you try to install a package and it fails with a message like "package xyzpackage needs xyzfile (not provided)" you can check for packages that include the missing file using the `provides` option as follows:

yum provides *missingfile*
With the `provides` option, yum will search your repositories for whatever file you enter (instead of *`missingfile`*) and return the name of any packages it finds that include that file.

Downloading Linux software
Instead of using yum to download and install software RPMS, you can simply browse around and download Linux software from the Internet using a Web browser (such as Mozilla) or an FTP program (such as the `ncftp` command). The browser often enables you to view the contents of an FTP site through a Web interface (look for an `index.html` file in an FTP directory). An `ftp` command (such as `ncftp`) has more options, but is less intuitive. There are also GUI-based FTP applications, such as `gFTP`, to make FTP services easier to use. (The `gFTP` command is described in Chapter 9.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The following procedures assume that you have a connection to the Internet.

Downloading with Mozilla

To download a Linux software package from the Internet using Mozilla, follow this procedure:

1. From the desktop panel, start Mozilla.
2. Type the name of an FTP site that has Linux software in the location box and press Enter.
3. To move around the FTP site, click Up to Higher Level Directory to move up, or click on a directory to move down.
4. When you find a package that you want to install, position the cursor over it, click the right mouse button, and then select Save Link Target As.
5. In the Save As window, click the Go Up a Level button to move up, or click a directory to go down until you find where you would like to save the package.
6. Click Save.

As the package is downloaded to your computer, a dialog box displays the progress. When the download is complete, the application is ready to be uncompressed and installed (or simply installed if you have an RPM file).

Downloading with ncFTP

If you want to use a text-based means of downloading files (instead of Mozilla), you can use any of several FTP commands that come with Red Hat Linux. One FTP client that I like to use is the `ncftp` command. (Other options are the `ftp` and `lftp` commands.) Here's an example of an `ncftp` procedure:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip The sftp command is a more secure way to connect to FTP servers. However, the FTP server needs to support ssh requests for sftp to work (which they don't all do).

Try an sft

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

1. From a shell or a Terminal window, type **`ncftp`** ***`location`***, where *`location`* is the name of an FTP site. For example:
2. $ **ncftp metalab.unc.edu**
   or
   $ **ncftp –u jake ftp.myveryownsrver.com**
   With no user name, as in the first example, `ncftp` logs you in as the anonymous user. With a user name (for example, `-u jake`), you are prompted to enter the password for that user at the FTP site.
3. When your login is accepted, you can use these commands to find the software package or document that you are looking for:
   ls — To list the contents of the current directory.
   cd *dir* — To change the current directory to the subdirectory *dir*. If you prefer, you can use two dots (`cd ..`) to go up a directory level. For example, try `cd /pub/Linux/apps/doctools.`
4. Type **binary** (to make sure the file is downloaded as a binary file).
5. To download a file from the current working directory, type **get *file*** where *file* is the application name. For example, to download the `whichman` application while `/pub/Linux/apps/doctools` is the current directory, type:
6. > **get whichman-2.0.tar.gz**
7. When the download is complete, type **exit**.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Before you start the ncftpcommand, make sure that your current directory is the one in which you want to download the file. Alternatively, you could change to the directory you want

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Understanding package names and formats

Whenever possible, you want to install the applications you use with Red Hat Linux from software packages in RPM format (files with a `.rpm` suffix). However, if an RPM isn't available, the software that you want may come in other package formats.

Say you just downloaded a file from the Internet that contains lots of names, numbers, dots, gzs, and tars. What does all that stuff mean? Well, when you break it down, it's really not that complicated.

Most of the names of archive files containing Linux applications follow the GNU-style package-naming conventions. The following example illustrates the package-naming format:

mycoolapp-4.2.3.i386.rpm

mycoolapp-4.2.3.tar.gz

mycoolapp-4.2.3.src.tar.gz

mycoolapp-4.2.3.bin.SPARC.tar.gz

mycoolapp-4.2.3.bin.ELF.static.tar.gz

These examples represent several different packages of the same software application. The name of this package is `mycoolapp`. Following the package name is a set of numbers that represent the version of the package. In this case, it is version 4.2.3 (the major version number is 4, followed by minor version number and patch level 2.3). After the version number is a dot, followed by some optional parts, which are followed by indications of how the file is archived and compressed.

The first line shows a package that is in Red Hat Package Manager (`.rpm`) format. The `.i386` before the `.rpm` indicates that the package contains binaries that are built to run Intel i386 architecture computers (in other words, PCs). See the sidebar "Using RPMs versus Building from Source" for the pros and cons of using prebuilt RPM binary packages as opposed to compiling the program yourself.

In the next two lines of the previous example, each file contains the source code for the package. The files that make up the package were archived using the `tar` command (`.tar`) and compressed using the `gzip` command (`.gz`). You use these two commands (or just the `tar` command with the `-z` option) to expand and uncompress the packages when you are ready to install the applications.

Between the version number and the `.tar.gz` suffixes there can be optional tags, separated by dots, which provide specific information about the contents of the package. In particular, if the package is a binary version, this information provides details about where the binaries will run. In the third line, the optional `.src` tag was added because the developer wanted to differentiate between the source and binary versions of this package. In the fourth line, the `.bin.SPARC` indicates that it is a binary package, ready to run on a SPARC workstation. The final line indicates that it is a binary package, consisting of statically linked ELF format executables.

Using Binary RPMs versus Building from Source

Binaries created in RPM format are easily installed, managed, and uninstalled using Red Hat tools. This is the recommended installation method for Red Hat Linux novices. Sometimes, however, building an application from source code may be preferable. Here are some arguments on both sides:

RPM — Installing applications from an RPM archive is easy. After the application is installed, there are both shell commands and GUIs for managing, verifying, updating, and removing the RPM application. You don't need to know anything about Makefiles or compilers. When you install an RPM package, RPM tools even check to make sure that other packages that the package depends on are installed. Because Red Hat has released RPM under the GPL, other Linux distributions also use it to distribute their software. Thus, most Linux applications are, or will be, available in RPM format.

Source code — Not all source-code packages are made into RPM binaries. If you use RPM, you may find yourself with software that is several versions old, when you could simply download the latest source code and run a few tar and make commands. Also, by modifying source code, you can tailor the package to better suit your needs.

For more information on RPMs, refer to the *Red Hat RPM Guide* by Eric Foster-Johnson (Red Hat

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Press/Wiley, 2003).

Here is a breakdown of the parts of a package name:

`name` — This is generally an all-lowercase string of characters that identifies the application.
dash (`-`)
version — This is shown as major to minor version number from left to right.
dot (`.`)
`src` or `bin` — This is optional, with `src` usually implied if no indication is given.
dot (`.`)
type of binary — This is optional and can include several different tags to describe the content of the binary archive. For example, i386 indicates binaries intended for Intel architectures (Pentium CPU) and SPARC indicates binaries for a Sparc CPU.
dot (`.`)
archive type — Often `tar` is used (`.tar`)
compression type — Often `gzip` is used (`.gzip`)

Using different archive and document formats

Many of the software packages that are not associated with a specific distribution (such as Red Hat or Debian) use the `tar/gzip` method for archiving and compressing files. However, you may notice files with different suffixes at software project sites.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Because we are using Red Hat Linux, most of the software applications we install are in Red Hat Package Management format (.rpm). Unless you want to build the software yourself—in

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 5-2 describes the different file formats that you will encounter as you look for software at a Linux FTP site. Table 5-3 lists some of the common document formats that are used in distributing information in Linux.

Table 5-2: Linux Archive File Formats

Format	Suffix	Description
Gzip file	`.gz` or `.z`	File was compressed using the GNU `gzip` utility. It can be uncompressed using the `gzip` or `gunzip` utilities (they are both the same).
Tar file	`.tar`	File was archived using the `tar` command. `tar` is used to gather multiple files into a single archive file. You can expand the archive into separate files using `tar` with different options.
Bzip2	`.bz2`	File was compressed with the `bzip2` program.
Tar/compressed	`.taz` or `.tz`	File was archived with `tar` and compressed with the UNIX `compress` command.
Linux Software Map	`.lsm`	File contains text that describes the content of an archive.
Debian Binary Package	`.deb`	File is a binary package used with the Debian Linux distribution. (See descriptions of how to convert Debian to Red Hat formats later in this chapter.)
Red Hat Package Management	`.rpm`	File is a binary package used with Red Hat Linux. Format also available to other Linux distributions.

Table 5-3: Linux Document Formats

Format	Suffix	Description
Hypertext Markup Language	`.html` or `.htm`	File is in hypertext format for reading by a Web browser program (such as Netscape Communicator).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 5-3: Linux Document Formats

Format	Suffix	Description
PostScript	`.ps`	File is in PostScript format for outputting on a PostScript printer.
SGML	`.sgml`	File is in SGML, a standard document format. SGML is often used to produce documents that can later be output to a variety of formats.
DVI	`.dvi`	File is in DVI, the output format of the LaTeX text-processing tools. Convert these files to PostScript or Hewlett-Packard's PCL using the `dvips` and `dvilj` commands.
Plain text		Files without a suffix are usually plain-text files (in ASCII format).

If you are not sure what format a file is in, use the `file` command as follows:

$ **file** ***archive-file***

This command tells you if it is a GNU `tar` file, RPM, `gzip`, or other file format. (This is a good technique if a file was renamed and lost its suffix.)

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Installing Red Hat Linux Applications

The first part of this section describes how to install and manage applications using tools created for RPM archive files. Later in this section, you will learn how to build and install applications that come as source code packages.

Note The rpm command can be used to install any packages that are in RPM format. If you need to install RPM packages that are part of your Red Hat Linux distribution, you can use the Package Management window to add and remove packages instead of the rpm command. The Package Management window is described in Chapter 2. To upgrade existing packages, you can use the up2date command, described in Chapter 10..

Installing and managing RPM files

When you get an application that is packaged in RPM format, you typically get a single file. The command used to work with RPM package files is `rpm`. To manage RPM packages, the `rpm` command has options that let you list all the packages that are installed, upgrade existing packages to newer versions, and query packages for information (such as the files or documentation included with the package). There is also a verify option to check that all files that make up the package are present and unchanged.

The `rpm` command has the following modes of operation:

- install (`-i`)
- upgrade (`-U`)
- freshen (-F)
- query (`-q`)
- verify (`-V`)
- signature check (`--checksig`)
- uninstall (`-e`)
- build (`-b`)
- rebuild database (`--rebuilddb`)
- fix permissions (`--setperms`)
- set owners/groups (`--setugids`)
- show RC (`--showrc`)

With these options, you can install RPM packages and verify that their contents are properly installed, correcting any problems that occur. You must be logged in as the root user to add or remove packages. You may, however, list installed packages, query packages for information, or verify a package's contents without root permission.

The following sections describe how to use `rpm` to install and work with your RPM applications.

Installing with rpm

To install an RPM archive file with the `rpm` command that is not yet installed on your system, most people generally use the same options they would if they were upgrading (the `-U` option). Here's an example of a command line you could use to install a new RPM package.

rpm -U [*options*] *package*

Package is the name of the RPM archive file. This package may be in the current directory, on a mounted CD (for example, /mnt/cdrom/RedHat/RPMS/*whatever*.rpm), or on an accessible FTP site (for example, ftp://sunsite.unc.edu/pub/Linux/games/strategy/galaxis-1.5-1.i386.rpm).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution Interrupting rpm during a package installation can leave stale lock files and possibly corrupt the database. As a result, subsequent rpm commands may hang. If this happens, you can probably correct the problem by removing old database locks. If that doesn't work, you can also try checking whether the database is corrupt and, if so, rebuilding the RPM database. Rebuilding the database can tak a long time, so only do it if the other options don't clear up the problem. Here's how to remove lock files, check the database, and rebuild the database (as root user):

```
# rm -f /var/lib/rpm/__db*
# db_verify /var/lib/rpm/Packages
# rpm --rebuilddb
```

Along with the `-U` option, you can use the following options to get feedback during a new installation:

-vv — Prints debugging information during installation. This is a good way to see everything that happens during the install process. (This output can be long, so you may want to pipe it to `more`.)

-h — Prints 50 hash marks (`#`) as the package unpacks. The intent is to see the progress of the unpacking process (so you can tell if the program is still working or stalled).

-percent — Prints the percentage of the total package that has been installed throughout the install process.

Before installing a package, `rpm` checks to make sure that it is not overwriting newer files or installing a package that has dependencies on other packages that are not installed. The following install options can be used to override conditions that may otherwise cause the installation to fail:

--force — Forces the contents of the current package to be installed, even if the current package is older than the one already installed, contains files placed there by other packages, or is already installed. (This is the same as using the `oldpackage`, `replacefiles`, and `replacepkgs` options.)

--oldpackage — Forces the package to be installed, even if the current package is older than the one already installed.

--replacefiles — Forces files in this package to be installed, even if the files were placed there by other packages.

--replacepkgs — Forces packages in this archive to be installed, even if they are already installed on the system.

--nodeps — Skips package dependency checks and installs the package, even if packages it depends on are not installed.

--ignorearch — Forces package to be installed, even if the binaries in the package don't match the architecture of your host computer.

--excludedocs — Excludes any man pages, texinfo documents, or other files marked as documentation.

--ignoreos — Forces package to be installed, even if the binaries in the package don't match the architecture of your operating system.

The following is a simple `rpm` command line used to install an RPM archive:

```
# rpm -U RealPlayer-8.0-1.i386.rpm
```

I personally like to see some feedback when I install something (by default, `rpm` is suspiciously quiet when it succeeds). Here is what the command looked like when I added the `-vv` option to get more verbose feedback, along with some of the output:

```
# rpm -Uvv RealPlayer-8.0-1.i386.rpm
D: ============== RealPlayer-8.0-1.i386.rpm
D: Expected size:      4978643 = lead(96)+sigs(100)+pad(4)+data(4978443)
D:   Actual size:      4978611
D: RealPlayer-8.0-1.i386.rpm: MD5 digest: OK (f8080e7c3c32eacc6912a5afb364479a)
D:      added binary package [0]
D: found 0 source and 1 binary packages
D: opening  db environment /var/lib/rpm/Packages joinenv
```

.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

.

D: installing binary packages

D: opening db environment /var/lib/rpm/Packages joinenv

D: opening db index /var/lib/rpm/Packages create mode=0x42

.

.

.

D: closed db index /var/lib/rpm/Basenames

D: closed db index /var/lib/rpm/Name

D: closed db index /var/lib/rpm/Packages

D: closed db environment /var/lib/rpm/Packages

From this output, you can see that `rpm` finds one binary package in this archive, verifies the checksum, opens the RPM database, installs the packages, and closes the database when done. Another way to verify that the install is actually working is to add the `-h` option, as follows:

rpm -Uvh RealPlayer-8.0-1.i386.rpm

RealPlayer ##

With the `-h` option, `rpm` chugs out 50 hash marks (#) until the package is done installing. As you can see, when everything goes well, installing with `rpm` is quite simple. Some problems can occur, however. Here are a couple of them:

Package dependencies errors — If the package you are installing requires an additional package for it to work properly, you will see an error noting the missing package. You should get and install that package before trying your package again. (You can override the failure with install options described above, but I don't recommend that because your package may not work without the dependent package.)

Nonroot user errors — If `rpm -U` is run by someone who is not the root user, the command will fail. The output will indicate that the `/var/lib/rpm` database could not be opened. Log in as root user and try again.

Upgrading packages with rpm

The upgrade option (`-U`) with `rpm` can, as you might expect, also be used to upgrade existing packages. The format is the same as described above:

rpm -U [*options*] *package*

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Although there is a separate install option (-i), I recommend using the -U option whether you are doing a new install or an upgrade. With -U, the package installs in either case. So rpm -U

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

One issue when upgrading is installing an older version of a package. For example, if you install a new version of some software and it doesn't work as well, you will want to go back to the old version. To do this, you can use the `--oldpackage` option as follows:

rpm -U --oldpackage AnotherLevel-0.7.4-1.noarch.rpm

If a later package of this name already exists, it is removed and the older version is installed.

Freshening packages with rpm

An option that is similar to the upgrade (`-U`) option is the freshen (`-F`) option. The main difference between the two is what happens if the RPM you are updating or freshening is not already installed on your Red Hat Linux system. With -U, the package is installed if it is not already installed; with `-F`, the package is not installed if it is not already installed.

A great use for freshen is when you have a directory full of updated RPM files that you want to install on your system. But, you only want to update those packages that are already installed. In other words, there may be a lot of RPMs in the directory you don't want. Freshen lets you just update the package you already have.

Let's say that you downloaded a directory of RPMs and you want to selectively freshen the ones you have installed. With the directory of RPMs as your current directory, you could type:

rpm -Fhv *.rpm

Packages already installed are updated with the new RPMs. All other RPMs are skipped.

Removing packages with rpm

If you no longer want to use a package (or you just want to recover some disk space), use the `-e` option to remove a package. In its simplest form, you use `rpm` with the `-e` option as follows:

rpm -e *package*

If there are no dependencies on this package, it is silently removed. Before you remove a package, however, you may want to do a quick check for dependencies. The `-q` option is used for a variety of query options. (Checking for dependencies isn't necessary because `rpm` checks for dependencies before it removes a package. You may want to do this for your own information, however.) To check for dependencies, do the following:

rpm -q --whatrequires *package*

If you decide to remove the package, I recommend using the `-vv` option with `rpm -e`. This lets you see the actual files that are being removed. I also suggest that you either direct the output to a file or pipe it to `more` because the output often runs off the screen. For example:

rpm -evv jpilot | more

This example removes the `jpilot` package and shows you the files that are being removed one page at a time. (Press the Spacebar to page through the output.)

Other options that you can run with `rpm -e` can be used to override conditions that would prevent the package from being removed or to prevent some processing (such as not running preuninstall and postuninstall scripts). Three of those options are as follows:

- **--nodeps** — Uninstall the package without checking for dependencies.
- **--noscripts** — Uninstall the package without running any preuninstall or postuninstall scripts.
- **--notriggers** — Uninstall the package without executing scripts that are triggered by removing the package.

If you feel nervous about boldly removing a package, you can always run the uninstall in test mode (`--test`) before you do the real uninstall. Test mode shows you everything that would happen in the uninstall without actually uninstalling. (Add the `--vv` option to see the details.) Here's an example:

rpm -evv --test jpilot | more

D: opening db environment /var/lib/rpm/Packages joinenv

D: opening db index /var/lib/rpm/Packages rdonly mode=0x0

D: locked db index /var/lib/rpm/Packages

D: opening db index /var/lib/rpm/Name rdonly mode=0x0

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

D: opening db index /var/lib/rpm/Pubkeys rdonly mode=0x0

.

.

.

D: closed db index /var/lib/rpm/Name

D: closed db index /var/lib/rpm/Packages

D: closed db environment /var/lib/rpm/Packages

If the results look fine, you can run the command again, without the `--test` option, to have the package removed.

Querying packages with rpm

Using the query options (`-q`) of `rpm`, you can get information about RPM packages. This can be simply listing the packages that are installed or printing detailed information about a package. Here is the basic format of an `rpm` query command (at least one option is required):

rpm -q [*options*]

The following are some useful options you can use with an `rpm` query:

-qa — Lists all installed packages.

-qf *file* — Lists the package that owns *file*. (The file must include the full path name or `rpm` assumes the current directory.)

-qi *package* — Lists lots of information about a package.

-qR *package* — Lists components (such as libraries and commands) that *package* depends on.

-ql *package* — Lists all the files contained in *package*.

-qd *package* — Lists all documentation files that come in *package*.

-qc *package* — Lists all configuration files that come in *package*.

To list all the packages installed on your computer, use the `-a` query option. Because this is a long list, you should either pipe the output to `more` or, possibly, use `grep` to find the package you want. The following command line displays a list of all installed RPM packages, and then shows only those names that include the string of characters `xfree`. (The `-i` option to `grep` says to ignore case.)

rpm -qa |grep -i xfree

If you are interested in details about a particular package, you can use the `rpm -i` query option. In the following example, information about the `dosfstools` package (for working with DOS file systems in Linux) is displayed:

rpm -qi dosfstools

Name : dosfstools Relocations: (not relocateable)

Version : 2.8 Vendor: Red Hat, Inc.

Release : 5 Build Date: Tue 18 Feb 2003 11:08:21 PM PST

Install Date: Sat 09 Aug 2003 02:15:35 PM PST

Build Host : daffy.perf.redhat.com

Group : Applications/System Source RPM:dosfstools-2.8-6.src.rpm

Size : 73278 License: GPL

Signature : (none)

Packager : Red Hat, Inc. <http://bugzilla.redhat.com/bugzilla>

Summary : Utilities for making and checking MS-DOS FAT

filesystems on Linux.

Description :

The dosfstools package includes the mkdosfs and dosfsck utilities,

which respectively make and check MS-DOS FAT filesystems on hard

drives or on floppies.

To find out about a package's contents, you can use the `-l` (list) option with your query. The following

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

example shows the complete path names of files contained in the `xpilot` package:

rpm -ql dosfstools | more

/sbin/dosfsck

/sbin/fsck.msdos

/sbin/fsck.vfat

/sbin/mkdosfs

/sbin/mkfs.msdos

/sbin/mkfs.vfat

/usr/share/man/man8/dosfsck.8.gz

.

.

.

Would you like to know how to use the components in a package? Using the `-d` option with a query will display the documentation (`man` pages, README files, HOWTOs, and so on) that is included with the package. If you are having trouble getting your X Window System running properly, you can use the following command line to find documents that may help:

rpm -qd XFree86 | more

/usr/X11R6/man/man1/XFree86.1x.gz

/usr/X11R6/man/man1/Xmark.1x.gz

/usr/X11R6/man/man1/Xserver.1x.gz

/usr/X11R6/man/man1/appres.1x.gz

/usr/X11R6/man/man1/atobm.1x.gz

/usr/X11R6/man/man1/bdftopcf.1x.gz

.

.

.

Many packages have configuration files associated with them. To see what configuration files are associated with a particular package, use the `-c` option with a query. For example, this is what you would type to find configuration files that are used with the `ppp` package:

rpm -qc ppp

/etc/pam.d/ppp

/etc/ppp/chap-secrets

/etc/ppp/options

/etc/ppp/pap-secrets

If you ever want to know which package a particular command or configuration file came from, you can use the `-qf` option. In the following example, the `-qf` option displays the fact that the `chgrp` command comes from the fileutils package:

rpm -qf /bin/chgrp

fileutils-4.1-4

Verifying packages with rpm

If something in a software package isn't working properly, or if you suspect that your system has been tampered with, the verify (`-V`) option of `rpm` can help you verify installed software against its original software package. Information about each installed package is stored on your computer in the RPM database. By using the verify option, you can check whether any changes were made to the components in the package.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The verify option uses the upper case letter (-V), while the verbose option uses the lower case (-v).

Various file size and permissions tests are done during a verify operation. If everything is fine, there is no output. Any components that have changed from when they were installed will be printed along with information indicating how they were changed. Here's an example:

rpm -V ppp

S.5....T c /etc/ppp/chap-secrets

S.5....T c /etc/ppp/options

S.5....T c /etc/ppp/pap-secrets

This output shows that the `ppp` package (used to dial-up a TCP/IP network such as the Internet) has had three files changed since it was installed. The notation at the beginning shows that the file size (`S`), the MD5 sum (`5`), and the modification time (`T`) have all changed. The letter `c` shows that these are all configuration files. By reviewing these files to see that the changes were only those that I made to get PPP working, I can verify that the software is okay.

The indicators that you may see when you verify the contents of a configuration file are:

- **5** — **MD5 Sum** — An MD5 checksum indicates a change to the file contents.
- **S** — **File size** — The number of characters in the file has changed.
- **L** — **Symlink** — The file has become a symbolic link to another file.
- **T** — **Mtime** — The modification time of the file has changed.
- **D** — **Device** — The file has become a device special file.
- **U** — **User** — The user name that owns the file has changed.
- **G** — **Group** — The group assigned to the file has changed.
- **M** — **Mode** — If the ownership or permission of the file changed.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip A utility is available to browse the contents of RPM files from Windows. With the rpmbrowser.exe utility, you can list and extract files from an RPM distribution. This utility is available

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Building and installing from source code

If no binary version of the package that you want is available, or if you just want to tailor a package to your needs, you can always install the package from source code. Source-code CDs for Red Hat Linux contain the source code equivalent (SRPMs) of the binary packages you installed. You can modify the source code and rebuild it to suit your needs.

Note To get CDs of the Red Hat Linux source code, see the mail-in coupon at the back of the book. Kernel source code is included on CD #1 with this book.

Software packages that are not available in RPM format are typically available in `tar/gzip` format. Although the exact instructions for installing an application from a source code archive vary, many

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

packages that are in the `.gz` and `.tar` formats follow the same basic procedure.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Before you install from source code, you will need to install a variety of software development packages. If you have the disk space, I recommend that you install all software development

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The following is a minimal list of C-programming software development tools:

gcc — Contains the `gcc` (GNU C compiler) compiler.
make — Contains the `make` command for making the binaries from Makefiles.
glibc — Contains important shared libraries, the C library, and the standard math library.
glibc-devel — Contains standard header files needed to create executables.
binutils — Contains utilities needed to compile programs (such as the assembler and linker).
kernel-source — Contains the Linux kernel source code and is needed to rebuild the kernel.
libc — Contains libraries needed for programs that were based on `libc` 5, so older applications can run on `glibc` (`libc` 6) systems.

Installing software in SRPMS format

To install a source package from CD #3 or the Red Hat Linux source CD, do the following:

1. Insert a Red Hat Linux source CD into the CD-ROM drive. It should mount automatically. (If it doesn't, type **`mount /mnt/cdrom`** in a Terminal window as root.)
2. Change to the source directory on the CD. For example:
3. # **cd /mnt/cdrom/SRPMS**
4. Choose the package you want to install (type **`ls`** to see the packages) and install it using the following command:
5. # **rpm -iv *packagename**.src.rpm**
   (Replace *packagename* with the name of the package you are installing.) The source is installed in the Red Hat Linux source tree (`/usr/src/redhat`). Spec files are copied to `/usr/src/redhat/SPECS`.
6. Change to the SPECS directory as follows:
7. # **cd /usr/src/redhat/SPECS**
8. Unpack the source code as follows:
9. # **rpm -bp *packagename**.spec**
10. The package's source code is installed to the `/usr/src/redhat/BUILD/`*`package`* directory, where *package* is the name of the software package.
11. You can now make changes to the files in the package's BUILD directory. Read the README, Makefile, and other documentation files for details on how to build the individual package.

Installing software in gzip/tar format

Here are some generic instructions that you can use to install many Linux software packages that are in `gzip/tar` format:

1. Get the source code package from the Internet or from a CD distribution and copy it into an empty directory (preferably using a name that identifies the package).
2. Assuming the file is compressed using `gzip`, uncompress the file using the following command:
3. # **gzip -d *package*.tar.gz**
   The result is that the package is uncompressed and the `.gz` is removed from the package name (for example, *`package`*`.tar`).
4. From the resulting `tar` archive, run the `tar` command as follows:
5. # **tar xvf *package*.tar**

   This command extracts the files from the archive and copies them to a subdirectory of the current directory.
6. Change directories to the new subdirectory created in Step 3, as follows:
7. # **cd *package***
8. Look for a file called `INSTALL` or `README`. One of these files should give you instructions on how to proceed with the installation. In general, the `make` command is used to install the package. Here are a few things to look for in the current directory:
   If there is a `Make.in` file, try running:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

./configure –prefix=/usr
make all

If there is an `Imake` file, try running:

xmkmf –a
make all

If there is a `Make` file, try running:

make all

After the program is built and installed, you may need to do additional configuration. You should consult the `man` pages and/or the HOWTOs that come with the software for information on how to proceed.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Even if you are not a programmer, reading the source code used to make a program can often give you insight into how that program works. Sometimes notes may be in the source

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To try out this procedure, I downloaded the `whichman` package, which includes utilities that let you find manual pages by entering keywords. The file I downloaded, `whichman-2.2.tar.gz`, was placed in a directory that I created called `/usr/sw/which`. I then ran the `gzip` and `tar` commands, using `whichman-2.2.tar.gz` and `whichman-2.2.tar` as options, respectively.
I changed to the new directory, `cd /usr/sw/which/whichman-2.2`. I then listed its contents. The README file contained information about the contents of the package and how to install it. As the README file suggested, I typed **`make`**, and then **`make install`**. The commands `whichman`, `ftwhich`, and `ftff` were installed in `/usr/bin`. (At this point, you can check the `man` page for each component to see what it does.)
The last thing I found in the README file was that a bit of configuration needed to be done. I added a `MANPATH` variable to my `$HOME/.bashrc` to identify the location of `man` pages on my computer to be searched by the `whichman` utility. The line I added looked like this:

```
export
MANPATH=/usr/share/man:/usr/man/man1:/usr/X11R6/man:/usr/share/doc/samba-
2.2.3a/docs
```

In case you are wondering, `whichman`, `ftwhich`, and `ftff` are commands that you can use to search for `man` pages. They can be used to find several locations of a `man` page, `man` pages that are close to the name you enter, or `man` pages that are located beneath a point in the directory structure, respectively.

Team LiB

PREVIOUS NEXT

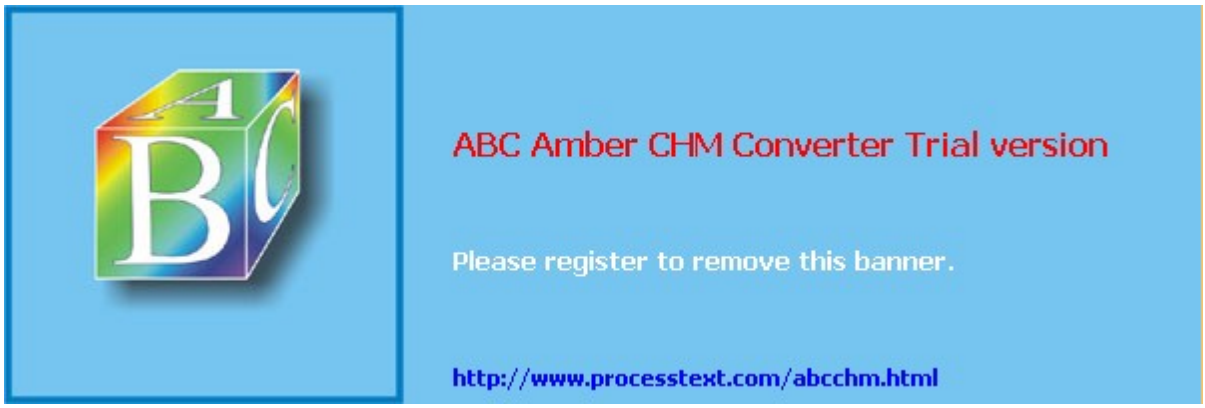

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Running X Window Applications

Setting up and configuring the X Window System to your liking is the hard part. By comparison, using X to run applications is relatively easy. If you have used Microsoft Windows operating systems, you already know the most basic ways of running an application from a graphical desktop. X, however, provides a much more flexible environment for running native Linux applications.

Cross-Reference See Chapter 3 for information on setting up an X desktop.

Starting applications from a menu

To run applications on your own desktop, most X window managers provide a menu, similar to the Microsoft Start menu, to display and select X applications. Applications are usually organized in categories. From the GNOME or KDE desktops in Red Hat Linux, open the Red Hat menu, select the category, and then select the application to run. Figure 5-1 shows an example of the Red Hat Main Menu and the Sound & Video subsubmenu.

Figure 5-1: Starting X applications from the Red Hat Menu.

Starting applications from a Run Program window

Not all installed applications appear on the menus provided with your window manager. For running other applications, some window managers provide a window, similar to the Run Program window, that lets you type in the name of the program you want to run.

To access the Run Program window:

1. Open the Red Hat Menu.
2. Click Run Program. The Run Program window appears.
3. Click Known Applications, click the program you want, and then click Run.
   If the application you want isn't on the list, you can either type the command you want to run (along with any options) and click Run, or you can click Append File to browse through directories to select a program to run. If you are running a program that needs to run in a Terminal window, such as the `vi` command, click the Run in Terminal button before running the command. Figure 5-2 is an example of the Run Program window.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

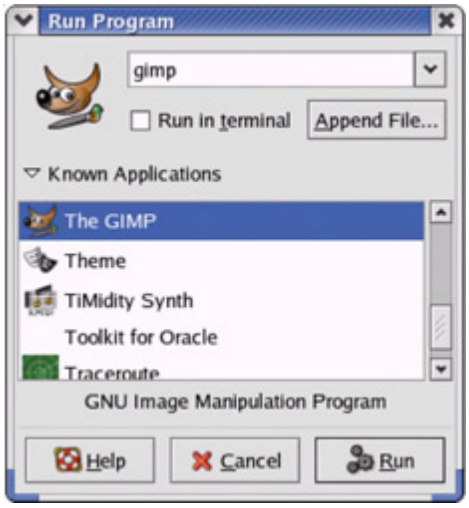

Figure 5-2: Select a program to run from the list in the Run Program window.

Starting applications from a Terminal window

I often prefer to run an X application, at least for the first time, from a Terminal window. There are several reasons why I prefer a Terminal window to selecting an application from a menu or Run Program window:

If there is a problem with the application, you see the error messages. Applications started from a menu or Run Program usually just fail silently.

Applications from menus run with set options. If you want to change those options, you have to change the configuration file that set up the menu and make the changes there.

If you want to try out a few different options with an application, a Terminal window is an easy way to start it, stop it, and change its options.

When you have found an application and figured out the options that you like, you can add it to a menu or a panel (if your window manager supports those features). In that way, you can run a program exactly as you want, instead of the way it is given to you on a menu.

Here is a procedure to run X applications from a Terminal window:

1. Open a Terminal window from your desktop (look for a Terminal icon on your Panel or a Terminal selection on a menu.)
2. Type
3. $ **echo $DISPLAY**

   The result should be something similar to the following:

   :0.0

   This indicates that the Terminal window will, by default, direct any X application you run from this window to display 0.0 on your local system. (If you don't see a value when you type that command, type **`export DISPLAY=:0.0`** to set the display value.)
4. Type the following command:
5. $ **xmms &**

   The `xmms` program should appear on your desktop, ready to work with. You should note the following:

   The `xmms` command runs in the background of the Terminal window (`&`). This means that you can continue to use the Terminal window while `xmms` is running. Because it is a child process of the Terminal, `xmms` will close if you close the Terminal window.

   I encountered no errors running `xmms` on this occasion. With other applications, however, text sometimes appeared in the Terminal window after the command was run. The text may say that the command can't find certain information or that certain fonts or colors cannot be displayed. That information would have been lost if the command were run from a menu.
6. If you want to know what options are available, type:
7. $ **xmms --help**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

8. Try it with a few options. For example, if you want to begin by playing a file and you have a WAV audio file named `file.wav`, you could type:
9. $ **xmms file.wav**
10. When you are ready to close the window, you can either do so from the `xmms` window (right-click on the `xmms` window and select Exit) or you can kill the process in the Terminal window. Type **`jobs`** to see the job number of the process. If it was job number 2, for example, you would type **`kill %2`** to kill the `xmms` program.

You should try running a few other X commands. Many X commands are stored in `/usr/X11R6/bin`.

Running remote X applications

X lets you start an application from anywhere on the network and have it show up on your X display. Instead of being limited by the size of your hard disk and the power of your CPU and RAM, you can draw on resources from any computer that gives you access to those resources.

Think about the possibilities. You can work with applications launched from any other computer that can run an X application — from a small PC to a supercomputer. Given the proper permission, you can work with files, printers, backup devices, removable drives, other users, and any other resources on the remote computer as though you were on that computer.

With this power, however, comes responsibility. You need to protect the access to your display, especially in networks where the other machines and users are not known or trusted. For example, you wouldn't want to allow anyone to display a login screen on your display, encouraging you to inadvertently give some cracker your login and password.

Traditionally, to run remote X applications, you basically only need to know how to identify remote X displays and how to use whatever security measures are put in place to protect your network resources. Using ssh to launch X applications is even simpler and more secure than the traditional method. Those issues are described in the following sections.

Traditional method to run remote X applications

If there is an X application installed on another computer on your network and you want to use it from your desktop, follow these steps:

Open permissions to your X server so that the remote application can use your display.
Identify your X server display to the application when it starts up.

When you run an X client on your local system, your local display is often identified as `:0`, which represents the first display on the local system. To identify that display to a remote system, however, you must add your computer's host name. For example, if your computer were named *whatever*, your display name would be:

whatever:0

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip In most cases, the host name is the TCP/IP name. For the computers on your local network, the name may be in your /etc/hosts file, or it may be determined using the Domain Name Sy

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You will probably use the display name in this form most of the time you run a remote X application. In certain cases, however, the information may be different. If your computer had multiple X displays (keyboard, mouse, and monitor), you may have numbers other than `:0` (`:1`, `:2`, and so on). It is also possible for one keyboard and mouse to be controlling more than one monitor, in which case you could add a screen number to the address, like this:

whatever:0.1

This address identifies the second screen (`.1`) on the first display (`:0`). The first screen is identified as `.0` (which is the default because most displays only have one screen). Unless you have multiple physical screens, however, you can skip the screen identifier.

There are two ways to identify your display name to a remote X application:

DISPLAY shell variable — The `DISPLAY` shell variable can be set to the system name and number identifying your display. After this is done, the output from any X application run from that shell will appear on the display indicated. For example, to set the `DISPLAY` variable to the first display on `whatever`, type one of the following:

export DISPLAY=whatever:0

or

setenv DISPLAY whatever:0

The first example shows how you would set the DISPLAY variable on a `bash` or `ksh` shell. The second example works for a `csh` shell.

-display option — Another way to identify a remote display is to add the `-display` option to the command line when you run the X application. This overrides the `DISPLAY` variable. For example, to open an `xterm` window on a remote system so that it appears on the first display on `whatever`, type the following:

xterm -display whatever:0

With this information, you should be able to run an X application from any computer that you can access from your local computer. The following sections describe how you may use this information to start a remote X application.

Launching a remote X application

Suppose you want to run an application from a computer named `remote1` on your local area network (in your same domain). Your local computer is `local1`, and the remote computer is `remote1`. The following steps show how to run an X application from `remote1` from your X display on `local1`.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

This procedure assumes that no special security procedures are implemented. It is the default situation and is designed for sharing applications among trusted computers (us

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

1. Open a Terminal window on the local computer.
2. Allow access for the remote computer (for example, `remote1`) to the local X display by typing the following from the Terminal window:
3.
4. **$ xhost +remote1**
   remote1 being added to access control list
5. Log in to the remote computer using any remote login command. For example:
6. $ **telnet –l *user* remote1**
   Password:
   Replace *user* with the name of the user login that you have on the remote computer. You will be prompted for a password.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note By default, the telnet service is not enabled in Red Hat Linux. The server (in this example, remote1) must consider the security consequences of enabling remote login services.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

7. Type the password for the remote user login. (You are now logged in as the remote user in the Terminal window.)
8. Set the `DISPLAY` variable on the remote computer to your local computer. For example, if your computer were named pine in the local domain, the command could appear as:
9. $ **export DISPLAY=pine:0**
   (If you are using a `csh` shell on the remote system, you may need to type **`setenv DISPLAY pine:0`**.)
10. At this point, any X application you run from the remote system from this shell will appear on the local display. For example, to run a remote Terminal window so that it appears locally, type:
11. $ **xterm -title "Terminal from Remote1"**
    The Terminal window appears on the local display with *Remote1* in the title bar.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip As a rule, I use the -title option when I run remote X applications. In the title, I try to indicate the remote computer name. That way, I am reminded that the application is not

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You need to remember some things about the remote application that appears on your display:

If you only use the login to run remote applications, you can add the line exporting the `DISPLAY` variable to a user configuration file on the remote system (such as `.bashrc`, if you use the `bash` shell). After that, any application that you run will be directed to your local display.

Even though the application looks as though it is running locally, all the work is being done on the remote system. For example, if you ran a word processing program remotely, it would use the remote CPU and when you save a file, it is saved to the remote file system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

Don't forget when a remote shell or file editor is open on your desktop. Sometimes people forget that a window is remote and will edit some important configuration file on the remote sy

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Securing remote X applications

There are several different traditional methods available to prevent just any X application from appearing on your display. Each method offers a different level of security, as well as different amounts of setup. The major security methods that are supported include:

Host Access — This method, described earlier as a default method, lets you decide which host computers can have access to your computer's X display. Use the `xhost` command to indicate those computers that are allowed access (`xhost +`*`hostname`*) and those that are not (`xhost -`*`hostname`*). If you allow a host access, any user from that computer is allowed access. This security method works best with small networks and single-user computers owned by people you trust.

Magic Cookie — This method, referred to as the MIT-MAGIC-COOKIE-1 method, lets you grant an individual user remote access to your display. The remote client must present a valid cookie before it can access your display.

The steps go something like this: You generate a magic cookie (using the `mcookie` command); you refer to the cookie when you start your X server (usually from a copy of the cookie in your `$HOME/.Xauthority` file); you distribute this cookie to remote computers that want to show X clients on your display; and the remote client refers to this cookie when it starts.

There are other X security methods that are also based on storing one or more randomly generated keys in the `.Xauthority` file. XDM-AUTHORIZATION-1 uses a two-part key — a DES encryption key (56-bit) and some random data (64-bit) — to enable access. When a system supports secure RPC, the SUN-DES-1 method can be used to enable access by setting up public and private key pairs that can be validated on a per user basis. Finally, the MIT-KERBEROS-5 method enables both the X client and X server to validate each other (using a Kerberos trusted third party).

Cross-Reference For more information about Kerberos, visit the Kerberos home page at web.mit.edu/kerberos/www/.

Using SSH to run remote X applications

Not only does the `ssh` command provide a secure mechanism for logging in to a remote system, it also provides a way of securely running remote X applications. After you log in to the remote computer using `ssh`, you can use that secure channel to forward X applications back to your local display. Here is an example:

1. Type the following `ssh` command to log in to a remote computer:
2. $ **ssh –l jake remote1**
3. jake@remote1's password: *******
4. Once you are logged in, type any X command and the window associated with that command appears on your local display. For example, to start the `gedit` command, type:
5. $ **gedit &**

For this to work, you don't need to open your local display (using `xhost`). The reason this works is because the SSH daemon (`sshd`) on the remote system sets up a secure channel to your computer for X applications. So as not to interfere with any real display numbers, the SSH daemon (by default) uses the display name of `localhost:10.0`.

This X forwarding feature is on by default with Red Hat Linux systems. The `ForwardX11 yes` option set in the `/etc/ssh/ssh_config` file (remotely) is what allows this to work.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Running Window, DOS, and Macintosh Applications

Linux is ready to run most applications that were created specifically for Linux, the X Window System, and many UNIX systems. Many other applications that were originally created for other operating systems have also been ported to Linux. There are still, however, lots of applications created for other operating systems for which there are no Linux versions.
Linux can run some applications that are intended for other operating systems using *emulator* programs. An emulator, as the name implies, tries to act like something it is not. In the case of an operating system, an emulator tries to present an environment that looks to the application like the intended operating system.

Note The most popular of these emulators, called WINE, is not really an emulator at all. WINE is a mechanism that implements Windows application-programming interfaces; rather than emulating Windows it provides the interfaces that a Windows application would expect. In fact, some people claim that WINE stands for "WINE Is Not an Emulator."

In the next sections, I discuss emulators that enable you to run applications that are intended for the following operating systems:

DOS
Windows 3.1
Windows 95
Windows 98
Windows 2000
Windows NT
Windows XP
Macintosh

Note In theory, any application that is Win32-compatible should be able to run using software such as WINE (described later). Whether or not a Windows application will run in an emulator in Linux must really be checked on a case-by-case basis.

Emulation programs for running DOS programs are available for Red Hat Linux. You need to install the `dosemu` and/or `xdosemu` packages (which are available from `rpmfind.net`). DOS emulation commands from these packages include: `dos`, `xtermdos`, `dosdebug`, `xdos`, and `dosexec`. Other operating system emulation programs can be added to your system, including:

WINE, which lets you run Windows 3.1, Windows 95, Windows 98, Windows 2000, Windows NT, and Windows XP binaries. (Windows NT and XP programs are not well supported.)
ARDI Executor, which enables you to run applications that are intended for the Macintosh operating system.

In general, the older and less complex the program, the better chance it has to run in an emulator. Character-based applications generally run better than graphics-based applications. Also, programs tend to run slower in emulation, due sometimes to additional debugging code put into the emulators.

Running DOS applications

Because Linux was originally developed on PCs, a variety of tools were developed to help developers and users bridge the gap between Linux and DOS systems. A set of Linux utilities called *mtools* lets you work with DOS files and directories within Linux. A DOS emulator called `dosemu` lets you run DOS applications within a DOS environment that is actually running in Linux (much the way a DOS window runs within a Windows operating system).

Using mtools

`Mtools` are mostly DOS commands that have the letter *m* in front of them and that run in Linux (though there are a few exceptions that are named differently). Using these commands, you can easily work with DOS files and file systems. Table 5-4 lists `mtools` that are available with Linux (if you have the `mtools` package installed).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 5-4: Mtools Available with Linux

Command	Function
`mattrib`	The DOS `attrib` command, which is used to change an MS-DOS file attribute flag.
`mbadblocks`	The DOS `badblocks` command, which tests a floppy disk and marks any bad blocks contained on the floppy in its FAT.
`mcd`	The DOS `cd` command, which is used to change the working directory to another DOS directory. (The default directory is A:\) that is used by other `mtools`.
`mcheck`	The DOS `check` command, which is used to verify a file.
`mcopy`	The DOS `copy` command, which is used to copy files from one location to another.
`mdel`	The DOS `del` command, which is used to delete files.
`mdeltree`	The DOS `deltree` command, which deletes an MS-DOS directory along with the files and subdirectories it contains.
`mdir`	The DOS `dir` command, which lists a directoyr's contents.
`mdu`	The Linux `du` command, which is used to show the amount of disk space used by a DOS directory.
`mformat`	The DOS `format` command, which is used to format a DOS floppy disk.
`minfo`	This command is used to print information about a DOS device, such as a floppy disk.
`mkmanifest`	This command is used to create a shell script that restores Linux filenames that were truncated by DOS commands.
`mlabel`	The DOS `label` command, which is used to make a DOS volume label.
`mmd`	The DOS `md` command, which is used to create a DOS directory.
`mmount`	This command is used to mount a DOS disk in Linux.
`mmove`	The DOS `move` command, which is used to move a file to another directory and/or rename it.
`mrd`	The DOS `rd` command, which is used to remove a DOS directory.
`mren`	The DOS `ren` command, which is used to rename a DOS file.
`mshowfat`	This command is used to show the FAT entry for a file in a DOS file system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 5-4: Mtools Available with Linux

Command	Function
`mtoolstest`	This command is used to test the `mtools` configuration files.
`mtype`	The DOS `type` command, which is used to display the contents of a DOS text file.
`mzip`	This command is used to perform operations with Zip disks, including eject, write protect, and query.

I usually use the `mtools` to copy files between my Linux system and a Windows system that is not on my network. I most often use `mcopy`, which lets me copy files using drive letters instead of device names. In other words, to copy the file `vi.exe` from floppy drive A: to the current directory in Linux, I would type:

mcopy a:\vi.exe

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

By default, the floppy-disk drive can be read from or written to only by the root user and the floppy group. To make the floppy drive accessible to everyone (assuming it is floppy drive A:),

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using dosemu

The DOS emulator package `dosemu` no longer comes with the Red Hat Linux distribution. To use `dosemu`, you can try to install source code available from `dosemu.sourceforge.net`.
With the utilities that come with the `dosemu` package, you can run DOS applications, as well as use your computer as if it were running DOS. This includes accessing hardware, working with DOS configuration files, and using the DOS file system.
The following commands can be used to start `dosemu` within Linux:

`dos` — Starts the DOS emulator in any shell.
`xdos` — Starts the DOS emulator in its own X window.
`xtermdos` — Starts the DOS emulator in an xterm window.
`dosdebug` — Starts a debug program to view information and error messages about a running DOS program.

Basic information about your DOS environment is set in the `/etc/dosemu.conf` file. The following list describes how the DOS environment is set up and how to change it:

CPU — The CPU is set to emulate an 80386 (Intel 386-compatible). You can change the `$_cpu` value to 80486 or 80586 (for Pentium).
Keyboard — The keyboard is set to `auto`, which tries to set the keyboard based on the current Linux console settings. You can specifically change the keyboard type by setting the `$_layout` value to a variety of other country/language combinations listed in the `dosemu.conf` file (such as `us` for US/English, `de` for German, or `it` for Italian).
X Settings — There are several X settings (beginning with `$_X`) that let you change things such as the title in the DOS window, the font used, and the cursor blink rate.
Floppy Disks — Floppy disk A (`/dev/fd0`) is set to a 3.5-inch floppy (`threeinch`) and floppy disk B (`/dev/fd1`) is not assigned. Either disk can be assigned to be 3.5-inch (`threeinch`), 5.25-inch (`fiveinch`), `atapi`, or `empty`.
Hard Disk Images — The file system that appears when you start DOS is actually a disk image file stored in `/var/lib/dosemu`. By default, `dosemu` uses `hdimage.first` (which is linked to `hdimage.freedos`). You can change that to another DOS image or to a DOS file system (such as `/dev/hda1`, if that partition were a DOS partition).
Serial Ports — No serial ports are assigned by default. You can assign any of the serial ports (`$_com1` to `$_com4`) to a device such as a modem (`/dev/modem`), a mouse (`/dev/mouse`), or a terminal line (`/dev/tty0`).

`Dosemu` is set up for the root user. It is not very secure to allow multiple users to have access to DOS because DOS does not have the same security protections for files and devices that Linux does. If you want other users on your Linux system to use DOS, however, edit the `/etc/dosemu.users` file so that it includes the following lines:

```
root c_all

all c_all
```

For more information on `dosemu`, visit the `dosemu` home page at `www.dosemu.org`.

Running Windows applications with WINE

The WINE project (`www.winehq.com`) has been making great strides in getting applications that were created for Microsoft Windows to run in Linux and other operating systems. WINE is not really an emulator, because it doesn't emulate the entire Windows operating system. Instead, because it implements Win32 application programming interfaces (APIs) and Windows 3.*x* interfaces, the WINE project is more of a "Windows compatibility layer."

To get WINE for your Red Hat Linux system, there are several places you can go:

Download sites — You can download WINE free of charge from several different Web sites. Check `winehq.com` for a list of download sites. I downloaded the latest RPMs built for Red Hat Linux from this site: `http://sourceforge.net/projects/wine/`. WINE daily builds (in RPM format) are available from `http://wine.dataparty.no`.
WineX — A commercial version of WINE, called WineX, is available from

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

TransGaming, Inc. (www.transgaming.com). TransGaming focuses on running Windows games in Linux, using WINE as its base. See Chapter 7 for descriptions of WineX.

CodeWeavers — If you need Microsoft Office or Web browser plug-ins, CodeWeavers (www.codeweavers.com) offers CrossOver Office and CrossOver Plug-in, respectively. Although these products cost some money, they offer friendly interfaces for installing and managing the Windows software and browser plug-ins that it supports.

While it's true that you can run many Windows applications using WINE, some fiddling is still required to get many Windows applications to work. If you are considering moving your desktop systems from Windows to Linux, the current state of WINE provides an opportunity to see if some Windows applications you need might run in Linux.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

I suggest that you test out any application you want to use in WINE before you put it on any computer that is a critical part of your work. Although I haven't done serious da

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Besides developing software, the WINE project maintains a database of applications that run under WINE. More than 1000 applications are listed, although many of them are only partially operational. The point is, however, that the list of applications is growing, and special attention is being paid to getting important Windows 98 and 2000 applications running.
Although not an open source products, VMware and Win4Lin are other good ways to run Windows applications along with a Linux system on the same running computer. With VMware you can simultaneously run Windows 98, Windows NT, Windows 2000, Windows XP, and Linux operating systems on the same PC. You can get more information about VMware from `www.vmware.com`. Win4Lin (`www.netraverse.com`) lets you run Windows applications in a running Linux system.
The next section describes how to set up Linux to run Windows applications using WINE.

Understanding WINE

For WINE to let you run Windows applications, it needs to have an environment set up that looks like a Windows system. When you install the wine package in Red Hat Linux, the package creates `/etc/wine/wine.conf` and `/etc/wine/system.reg` files that, like the Windows registry, identify the locations of components an application would need in a Windows operating system.
The location of the basic Windows operating-system directories for WINE is the `/usr/share/wine-c` directory (in other words, drive C:) by default. Table 5-5 shows how C: and other Windows drive letters are mapped into Linux directories and devices.

Table 5-5: Windows Drive Letters in WINE

Description	Drive Letter or Name	Linux Directory or Output Location	Linux Device Name (If Applicable)
Floppy drive	A:	`/mnt/floppy`	/dev/fd0
Hard disk #1	C:	`/usr/share/wine-c`	
CD-ROM	D:	`/mnt/cdrom`	/dev/cdrom
Temp directory	E:	`/tmp`	
User's home directory	F:	`$HOME` (user's home directory)	
Root directory	Z:	`/`	
Serial ports	COM1 COM2 COM3 COM4		/dev/ttyS0 /dev/ttyS1 /dev/ttyS2 /dev/ttyS3
Parall ports	Lpt1		/dev/lp0
Print spooler	LPT1	Pipes printer jobs to the `lpr` command `(| lpr)`	
WineLook	Win95 (how the system appears to apps: Win31, Win98, or the default Win95)		

Within the `/usr/share/wine-c` directory (that is, your C: drive), you should see some things that are familiar to you if you are coming from a Windows environment: `autoexec.bat` and `config.sys`

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

files (empty to start with), and `My Documents`, `Program Files`, and `windows` directories. Assuming that the basic directory structure is okay with you, you may still want to consider changing a few items. For example, you may want to change the command to print documents (set in the `/etc/wine.conf` file) to something other than the "`| lpr`" command to use a particular printer or add some Linux printing options.

Setting up Windows applications

The best way to understand how to set up and run a Windows application is to step through the process of getting a simple Windows accessory to run in WINE. To do this procedure, you will need the following:

The wine package installed in Red Hat Linux.
Access to MS Windows components, including the executable program you want to run and the DLL files the executable needs to run.

If you have a dual-boot system (Red Hat Linux and Windows), the best way to do this procedure is to have Linux running with the Windows partition mounted somewhere so that you can copy the files you need. Otherwise, you should skip the first step and copy the files you need from a floppy or CD, or over the network.
In this example, we want to be able to run the Microsoft Paint program (`mspaint.exe`) from Windows in Linux using WINE. Here is an example of how to go about it.

1. If you have a dual-boot system start by mounting the Windows partition. Here's how:

   As the root user, from a Terminal window, type:
   # **fdisk –l**
   This lists your hard-disk partitions. Look for a partition that is listed as something like WIN95 FAT32.
   Make a directory on which to mount your Windows partition. For example, to use the `/mnt/win` directory, type:
   # **mkdir /mnt/win**
   Assuming that the Windows partition is on the first partition of the first IDE hard disk (`/dev/hda1`), type the following to mount the Windows partition:
   # **mount –t vfat /dev/hda1 /mnt/win**
   At this point you should have a mounted Windows partition available from `/mnt/win`.
   To make that partition permanently available to you from Linux, add the following line to your `/etc/fstab` file:
   /dev/hda1 /mnt/win vfat defaults 0 0
2. The preceding entry causes the Windows partition to mount automatically on `/mnt/win` each time your Red Hat Linux system boots up.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note For the purposes of the rest of this procedure, I assume that you have a Windows system accessible from /mnt/win. If this is not the case, you need to find a way to copy file

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

3. Copy the `mspaint.exe` program to a location accessible from your Linux computer. For example, with a mounted Windows partition as previously described, you could type:
4. # **cp /mnt/win/Program\ Files/Accessories/mspaint.exe /bin**
5. Run the `wine` command with `mspaint.exe` as an argument to see if it can run or if it needs some added DLL files.
6. # **wine /bin/mspaint.exe**
7. err: module:PE_fixup_imports Module (file) MFC42.DLL (which is needed by Z:\a\mspaint.exe) **not** found
   This says that you need the `mfc42.dll` file to be able to run `mspaint.exe`.
8. Next, copy the necessary DLL file from the Windows partition. For example:
9. # **cd /usr/share/wine-c/windows/system**
10. # **cp /mnt/win/windows/system/mfc42.dll mfc42.dll**
11. With the proper DLL files installed, type **`wine /bin/mspaint.exe`** again and the MS Paint window opens. An example of the Paint window running in Red Hat Linux is shown in Figure 5-3.

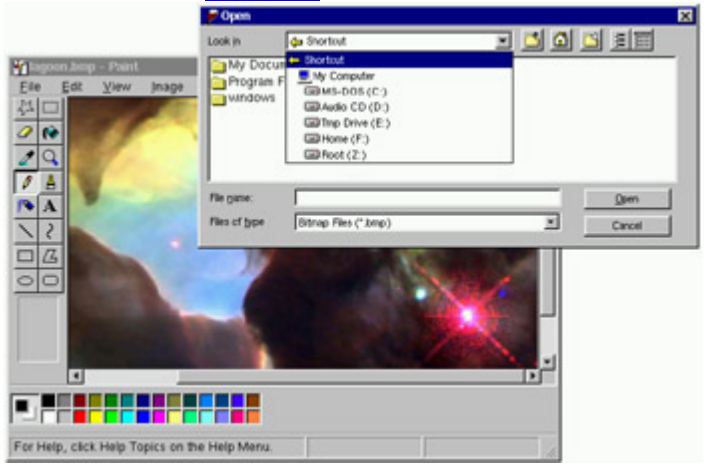

Figure 5-3: Running Paint in Red Hat Linux using WINE.

In Figure 5-3, Paint is displaying a bitmap image that I opened from the Red Hat Linux file system. Choose File → Open. In the Open dialog box, you can see how the application views the locations available to it. The C:, D:, E:, F:, and Z: drives all represent the locations shown in Table 5-5. (The computer in the example didn't have a floppy, or you would see an A: too.)

Finding Windows applications for WINE

Applications that have been tested to run under WINE are maintained at the Wine Application Database (`http://appdb.codeweavers.com`/). The database is divided into the following categories:

Educational Software and Computer-Based Training
Games
Multimedia
Networking & Communications
Productivity
Programming / Software Engineering
Reference, Documentation, and Information
Scientific, Technical, and Mathmatical
Special Purpose
Utilities

Hundreds of applications have been submitted for the database. Most of the packages that have been tested are games. In most cases, the person listing the application describes the platform he or she tested on and any experiences he or she had getting it to run. The host of the list, `codeweavers.com`, asks that you create a login to access the database.

Running Macintosh applications with ARDI Executor

Besides enabling you to run many popular Macintosh applications on the PC, ARDI Executor (`www.ardi.com`) from ARDI, Inc. lets you work with Mac-formatted floppies and a variety of Mac drives. Find more about ARDI Executor from `www.ardi.com/executor.php`.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ARDI also maintains a listing of compatible Mac software in its Compatibility Database. There are literally hundreds of Mac applications listed. Each application is color coded (green, yellow, orange, red, or black) to indicate how well the software runs under ARDI. Green and yellow are fully usable and largely usable, respectively. Orange is mostly unusable. Red means the application won't run at all, and black means it won't run because it requires features that aren't implemented in ARDI. At a glance, about two-thirds of the applications listed were either green or yellow.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

Between applications written directly for Linux and other UNIX systems, those that have been ported to Linux, and those that can run in emulation, hundreds of applications are available to be used with Red Hat Linux. There are dozens of locations on the Internet for downloading Linux applications, and many more that can be purchased on CD.
To simplify the process of installing and managing your Linux applications, Red Hat developed the Red Hat Package Manager (RPM). Using tools developed for RPM, such as the `rpm` command, you can easily install, remove, and perform queries on Linux RPM packages.

Of the types of applications that can run in Linux, those created for the X Window System provide the greatest level of compatibility and flexibility when used in Linux. However, using emulation software, it is possible to run applications intended for DOS, Windows 95/98/2000/NT/XP, and Macintosh operating systems.

Team LiB

PREVIOUS NEXT

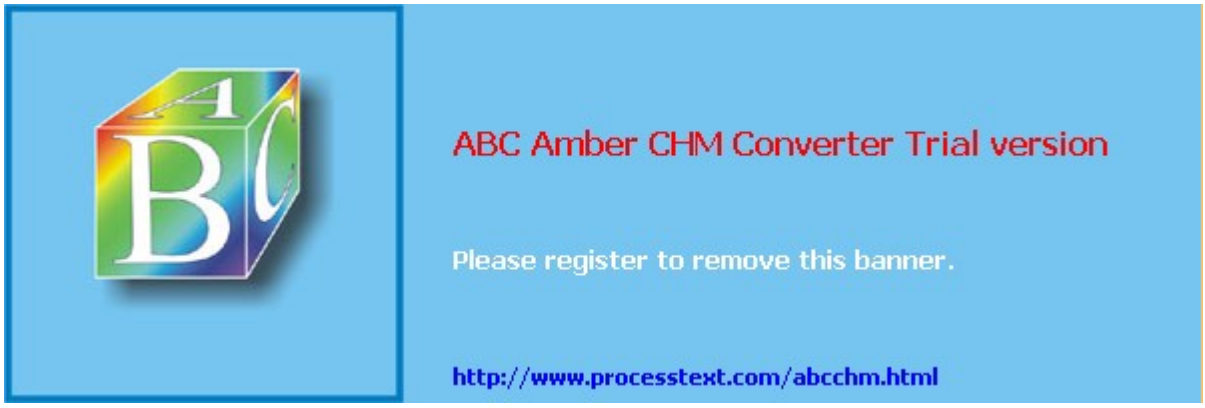

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 6: Publishing with Red Hat Linux

Overview

To survive as a desktop system, an operating system must be able to perform at least one task well: produce documents. It's no accident that, after Windows, Word (now bundled into Microsoft Office) is the foundation of Microsoft's success on the desktop. Red Hat Linux includes tools for producing documents, manipulating images, scanning, and printing. Almost everything you would expect a publishing system to do, you can do with Red Hat Linux.

OpenOffice is a powerful open-source office suite available for the first time as part of the Red Hat Linux distribution. Based on Sun Microsystem's Star Office, OpenOffice includes a word processor, spreadsheet, presentation manager, and other personal productivity tools.

The first document and graphics tools for Linux were mostly built on older, text-based tools. Recently, more sophisticated tools for writing, formatting pages, and integrating graphics have been added. Despite their age, many of the older publishing tools (such as Groff and LaTeX) are still used by people in the technical community.

In this chapter, I describe both text-based and GUI-based document preparation software for Red Hat Linux. I also describe tools for printing and displaying documents, as well as software for working with images.

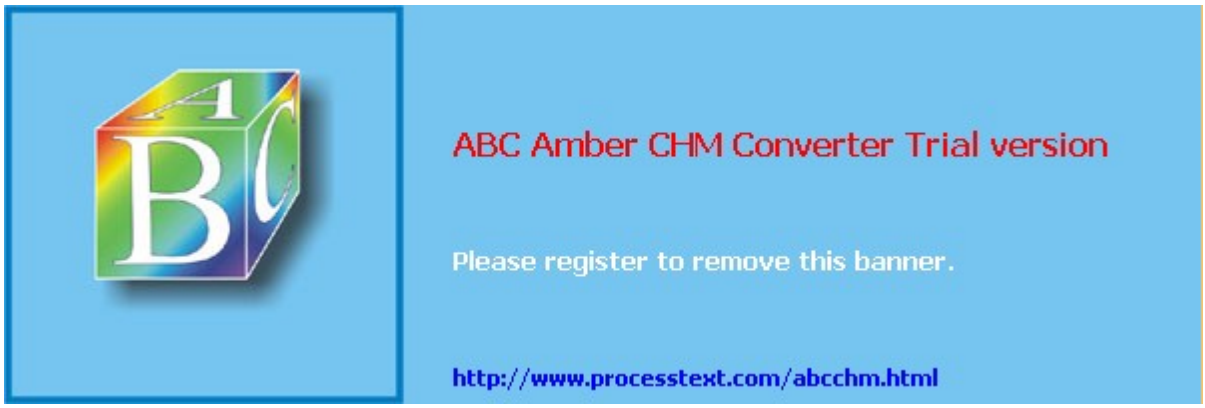

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using OpenOffice

Some have called OpenOffice a significant threat to Microsoft's dominance of the desktop market. If a need to work with documents in Microsoft Word format has kept you from using Linux as your desktop computer, OpenOffice is a big step toward removing that obstacle.

Note If you are willing to pay a few dollars, CrossOver Office from Codeweavers.com lets you install and run older versions of Microsoft Office (97, 2000, XP, and so on) from your Linux desktop. See Chapter 5 for further information or check out www.codeweavers.com/products/office.

Red Hat Linux includes the entire OpenOffice suite of desktop applications. Based on source code from Sun Microsystem's StarOffice, OpenOffice consists of the following office-productivity applications:

OpenOffice Writer — A word-processing application that can work with documents in file formats from Microsoft Word, StarOffice, and several others. Writer also has a full set of features for using templates, working with fonts, navigating your documents, including images and effects, and doing tables of contents.

OpenOffice Calc — A spreadsheet application that lets you incorporate data from Microsoft Excel, StarOffice, Dbase, and several other spreadsheet formats. Some nice features in Calc let you create charts, set up database ranges (to easily sort data in an area of a spreadsheet), and use data pilot to arrange data in different points of view.

OpenOffice Draw — A drawing application that lets you add, edit, and align objects; incorporate textures; include textures and colors; and work with layers of objects. It lets you incorporate images, vector graphics, AutoCAD, and a variety of other file formats into your drawings. Then, you can save your drawing in OpenOffice Drawing or StarDraw formats.

OpenOffice Math — A calculation program that lets you create mathematical formulas.

OpenOffice Impress — A presentation application that includes several types of slide effects. Using Impress, you can create and save presentations in Microsoft PowerPoint, StarDraw, and StarImpress formats.

Unlike other applications that were created to work with Microsoft document and data formats, OpenOffice (although not perfect) does a very good job of opening and saving those files with fewer problems. Basic styles and formatting that open in OpenOffice often don't look noticeably different from the way they appear in Microsoft Office.

Icons for launching Writer, Impress, and Calc are right on the Red Hat desktop panel (just click to start one). Alternatively, to open those and other OpenOffice applications, click Office from the main menu. Then select the OpenOffice application you want to open. Figure 6-1 shows an example of OpenOffice Writer displaying and working with a document file that was originally created by Microsoft Word.

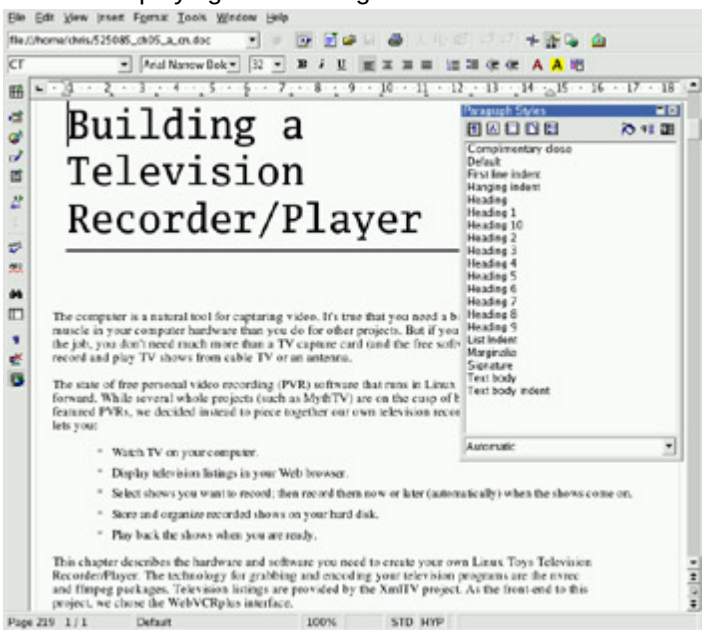

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Figure 6-1: Work with Microsoft Word documents in OpenOffice Writer.

The controls in OpenOffice are the same ones you would find in Microsoft Word. Toolbars include boxes for identifying filenames, changing styles, font types, and font sizes. Buttons let you save and print the file, change justification, and cut, copy, and paste text. In other words, Writer includes almost everything you expect in an advanced word processor.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Other Word Processors

Other free word processors that come with Red Hat Linux include AbiWord and KOffice (which is part of the KDE desktop). As for commercial offerings, there is StarOffice from Sun Microsystems.

StarOffice — The StarOffice office productivity suite contains applications for word processing, spreadsheets, presentation graphics, e-mail, news, charting, and graphics. It was created to run on Linux systems, but it runs in other environments as well. It can import and export a variety of Microsoft file formats. StarOffice is owned by Sun Microsystems, which sells it as a commercial product.

AbiWord — The AbiWord word processor (`abiword` command) is the first application produced by the AbiSource project (`www.abisource.com`). Besides working with files in its own AbiWord format (.abw and .zabw), AbiWord can import files in Microsoft Word and several other formats. If AbiWord is not installed, you can install the abiword package from the third installation CD (CD #3) that comes with this book.

KOffice — The KOffice package contains a set of office productivity applications designed for the KDE desktop. It includes a word processor (KWord), spreadsheet (KSpread), a presentation creator (KPresenter), and a diagram drawing program (KChart). These applications can be run separately or within a KOffice Workspace. (The koffice package is on CD #3 of the disks that come with this book.)

Using StarOffice

The StarOffice suite from Sun Microsystems Inc. (`www.sun.com/staroffice`) is a product that runs on Linux, UNIX, and Windows operating systems. Like OpenOffice, StarOffice contains many features that make it compatible with Microsoft Office applications. In particular, it includes the capability to import Microsoft Word and Excel files. There is also strong support for producing HTML files.

StarOffice is probably the most complete integrated office suite for Linux. It includes:

StarWriter — This is the StarOffice word processing application. It can import documents from a variety of formats, with special emphasis on Word documents.

StarCalc — This is the spreadsheet program that comes with StarOffice. You can import spreadsheets from Microsoft Excel and other popular programs.

StarImpress — You can create presentations with the StarImpress application.

StarDraw — This is a vector-oriented drawing program. It includes the capability to create 3D objects and to use texturing.

StarBase — You can manage your data with StarBase. It can access a variety of database interfaces.

StarSchedule — Use this application to manage your schedule, track events, and store your contacts. You can even interface StarSchedule to a PalmPilot.

There are also other applications in StarOffice that create business graphics (StarChart), edit images (StarImage), edit mathematical formulas (StarMath), and manage e-mail (StarMail). You can manage all these applications from the StarDesktop. The StarDesktop features an explorer tree for navigation and a task bar (so you can easily switch between tasks).

You can download StarOffice 6.0 for Linux or purchase a boxed set from the StarOffice Web site at `www.sun.com/staroffice`. Although StarOffice was once available free for download, the current price to download the software for home users is $75.95.

AbiWord

The AbiWord word processor is a very nice, free word processor from the AbiSource project (`www.abisource.com`). If you are creating documents from scratch, AbiWord includes many of the basic functions you need to create good-quality documents.

With AbiWord, you can select what type of document the file contains. You can select to read the file in the following formats:

AbiWord (.abw)
GZipped AbiWord (.zabw)
Rich Text Format (.rtf)
Microsoft Word (.doc)
UTF8 (.utf8)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Text (.txt)

AbiWord doesn't yet import all of these file types cleanly. Although MS Word styles have recently become supported in AbiWord, sometimes tables, graphics, and other features don't translate perfectly. If you want to work with a Microsoft Word document in AbiWord, open it as AbiWord, correct any font problems, and save the document in AbiWord format. AbiWord has vastly improved in the past few releases, but you can still have problems if you need to exchange files with others who are using Microsoft Word. (To keep files in Word format, OpenOffice and StarOffice work much better.)

AbiWord is a great first try as a usable word processor. Recently added features, such as styles and bullets, continue to make AbiWord a more useful publishing tool. It's not competitive with comparable commercial products, but its developers continue to improve it.

Using KOffice

There is now a KDE office suite of applications that go with the KDE desktop. The KOffice package has the basic applications you would expect in an integrated office suite: a word processor (KWord), spreadsheet (KSpread), a presentation creator (KPresenter), and a diagram drawing program (KChart).

You can start by opening the KOffice Workspace (from the KDE or Red Hat menu, Office → More Office Applications → KOffice Workspace). From the workspace window that opens, you can select from the different office applications from the left column. Open multiple documents in any of the applications. Then click on Documents in the left column to choose which one to display at the moment. Figure 6-2 shows the KOffice workspace, displaying a KWord document.

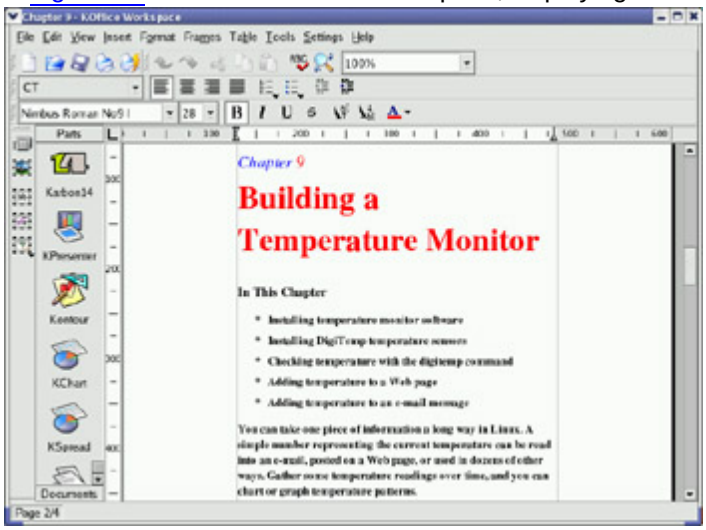

Figure 6-2: The KOffice Workspace lets you work with multiple KDE office applications at once.

You can work with a variety of document, spreadsheet, and image types; not many commercial document types are supported yet. So you may need to import documents using other tools before you can read them into KWord. The KSpread can open several different spreadsheet styles, however, such as Microsoft Excel and GNUmeric spreadsheets.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using Traditional Linux Publishing Tools

With old-school text processors (such as Groff and TeX), you can ignore document appearance while writing. Plain-text macros tell post-processors how to lay out a document for printing after writing is done. With word processors (such as OpenOffice Word and StarOffice), you mark up text and see the basic layout of the document as you write.

Some attributes of the traditional Linux publishing tools make them particularly well-suited for certain types of document publishing. Groff and LaTeX (which is based on TeX) come with Red Hat Linux and have been popular among technical people. Reasons for that include:

- You can manipulate files in plain text. Using tools such as `sed` and `grep`, you can scan and change one document or hundreds with a single command or script.
- Scientific notation is supported. With `geqn`, you can create complex equations. LaTeX and TeX are suited for technical notation. Some math publications require LaTeX.
- Editing can be faster because traditional Linux documents are created with a text editor. You usually get better performance out of a text editor than a word processor.

Simple page layouts work well with Linux documentation tools. For example, a technical book with a few flow charts and images can be easily produced and maintained using Groff or TeX documentation tools. Letters and memos are also easy to do with these tools. And, of course, Linux man pages are created with text-based tools.
Also, Linux likes PostScript. Although people think of PostScript as a printing language, it is really more of a programming language (you could write PostScript code directly). Most Linux document-processing software includes print drivers for PostScript. Also, some documents on the Web are distributed in PostScript (`.ps`).

The drawback to the traditional Linux document tools is that they are not intuitive. Although there are some easier front-ends to LaTeX (see the description of LyX later on), if you are creating documents in a text editor, you need to learn what macros to type into your documents and which formatting and print commands to use.

Note For many years, the UNIX system documentation distributed by AT&T was created in troff/nroff formats, which predate Groff. The documents used separate macro packages for man pages and guide material. Using a source code control system (SCCS), thousands of pages of documentation could be ported to different UNIX systems. Today, Red Hat still includes the same tools to work with man pages.

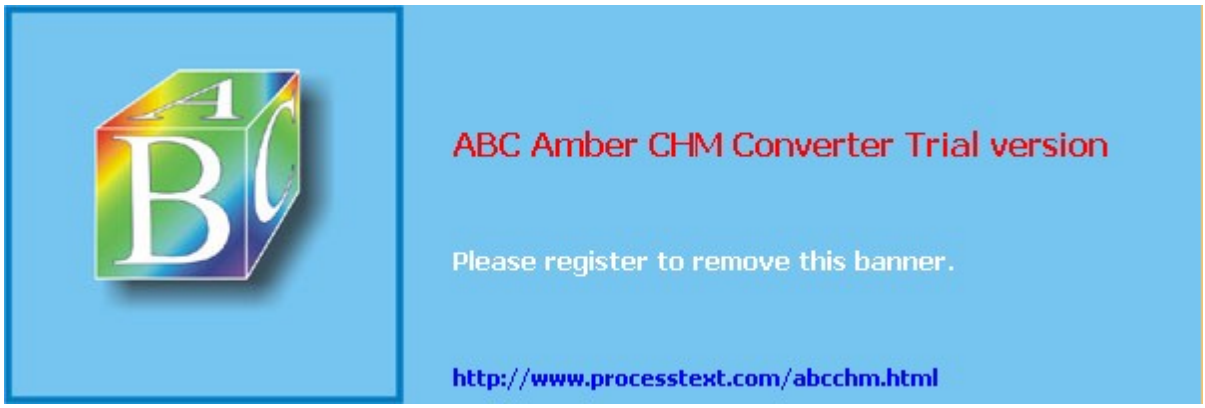

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Creating Documents in Groff or LaTeX

You can create documents for either of Linux's Groff (troff/nroff) or LaTeX (TeX) styles of publishing using any text editor. Red Hat Linux comes with several text editors, or you can download others from the Internet. See the "Choosing a Text Editor" sidebar for information.

The process of creating documents in Groff or LaTeX consists of the following general steps:

1. Create a document with any text editor. The document will contain text and markup.
2. Format the document using a formatting command that matches the style of the document that you created (for example, with `groff` or `latex`). During this step, you may need to indicate that the document contains special content, such as equations (`eqn` command), tables (`tbl` command), or line drawings (`pic` command).
3. Send the document to an output device. The device may be a printer or display program.

If you are used to a word processor with a GUI, you may find these publishing tools difficult. In general, Groff is useful to create man pages for Red Hat Linux. LaTeX is useful if you need to produce mathematical documents, perhaps for publication in a technical journal.

Text processing with Groff

The `nroff` and `troff` text formatting commands were the first interfaces available for producing typeset quality documents with the UNIX system. They aren't editors; rather, they are commands that you send your text through, with the result being formatted pages:

`nroff` (Produces formatted plain text and includes the ability to do pagination, indents, and text justification, as well as other features.

`troff` (Produces typeset text, including everything `nroff` can do, plus the ability to produce different fonts and spacing The `troff` command also supports kerning.

[redacted]

Choosing a Text Editor

Hardc[redacted]ux users tend to edit files with either the vi or emacs text editor. These editors have been [redacted] time and are hard to learn, but efficient to use. (Your fingers never leave the keyboard.) The emacs editor has some GUI support, though it will run fine in a Terminal window. There are also GUI versions of vi and emacs that add menu and mouse features to the editors. These are GVim (`gvim` command in the vim-X11 package) and Xemacs (`xemacs` command) editors.

Some of the other, simpler text editors that can run on your graphical desktop are:

gedit (`gedit` command) — This text editor, which comes with Red Hat, is the lightweight text editor for GNOME. It has simple edit functions (cut, copy, paste, and select all) and settings let you set indentations and word wrap. Special functions, such as a spell checker and a diff feature are included. You can start `gedit` by typing `gedit` from a Terminal window. Go to `http://gedit.sourceforge.net` for more information.

Advanced Editor (`kwrite` command) — This text editor includes a menu bar to create, open, or save files. It also has simple edit functions (cut, copy, paste, undo, and help). Other edit features let you set indents, find/replace text, and select all text. This comes with the KDE desktop, so you can access it by selecting Accessories → More Accessories → Kwrite.

Text Editor (`kedit` command) — Another simple text editor. Features let you open files from your file system or from a URL. It also includes a convenient toolbar and a spell checker. It comes with the KDE desktop, so you can access it by selecting Accessories → More Accessories → Text Editor.

[redacted]

The `groff` [redacted]he front-end for producing `nroff/troff` documentation. Because Linux man pages are fo[redacted]output in Groff, most of the examples here help you create and print man

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

pages with Groff.

People rarely use primitive `nroff/troff` markup. Instead, there are common macro packages that simplify creating `nroff/troff` formatted documents, which include:

man — The `man` macros are used to create Linux man pages. You can format a man page using the `-man` option to the `groff` command.

mm — The `mm` macros (memorandum macros) were created to produce memos, letters, and technical white papers. This macro package includes macros for creating a table of contents, lists of figures, references, and other features that are helpful for producing technical documents. You can format an `mm` document using the `-mm` `groff` option.

me — The `me` macros were popular for producing memos and technical papers on Berkeley UNIX systems. Format an `me` document using the `-me` `groff` option.

Groff macro packages are stored in `/usr/share/groff/*/tmac`. The `man` macros are called from the `an.tmac` file, `mm` macros are from `m.tmac`, and `me` macros are from `e.tmac`. The naming convention for each macro package is *xxx*`.tmac`, where the *xxx* is replaced by one or more letters representing the macro package. In each case, you can understand the name of the macro package by adding an `m` to the beginning of the file suffix.

Tip Instead of noting a specific macro package, you can use -mandoc to choose a macro package.

When you run the `groff` formatting command, you can indicate on the command line which macro packages you are using. You can also indicate that the document should be run through any of the following commands that preprocess text for special formats:

eqn — This preprocessor formats macros that produce equations in `groff`.

pic — This preprocessor formats macros that create simple line drawings in `groff`.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

tbl — This preprocessor formats macros that produce tables within `groff`.

The formatted Groff document is output for a particular device type. The device can be a printer, a window, or (for plain text) your shell. Here are output forms supported by Groff:

ps — Produces PostScript output for PostScript printer or a PostScript previewer.
lj4 — Produces output for an HP LaserJet4 printer or other PCL5-compatible printer.
ascii — Produces plain-text output that can be viewed from a Terminal window.
dvi — Produces output in TeX dvi, to output to a variety of devices described later.
X75 — Produces output for an X11 75 dots/inch previewer.
X100 — Produces output for an X11 100 dots/inch previewer.
latin1 — Produces typewriter-like output using the ISO Latin-1 character set.

Formatting and printing documents with Groff

You can try formatting and printing an existing Groff document using any man pages on your Red Hat Linux system (such as those in `/usr/share/man/*`). (Those man pages are compressed, so you can copy them to a temporary directory and unzip them to try out Groff.)
These commands copy the `chown` man page to the `/tmp` directory and unzips it. Then, `groff` formats the `chown` man page in plain text so you can page through it on your screen.

$ **cp /usr/share/man/man1/chown.1.gz /tmp**

$ **gunzip /tmp/chown.1.gz**

$ **groff -Tascii -man /tmp/chown.1 | less**
In the previous example, the `chown` man page (`chown.1.gz`) is copied to the `/tmp` directory, is unzipped (using `gunzip`), and is output in plain text (`-Tascii`) using the man macros (`-man`). The output is piped to `less`, to page through it on your screen. Instead of piping to `less` (`| less`), you could direct the output to a file (`> /tmp/chown.txt`).

To format a man page for typesetting, you could specify PostScript or HP LaserJet output. You should either direct the output to a file or to a printer. Here are a couple of examples:

$ **groff -Tps -man /tmp/chown.1 > /tmp/chown.ps**

$ **groff -Tlj4 -man -l /tmp/chown.1**
The first example creates PostScript output (`-Tps`) and directs it to a file called `/tmp/chown.ps`. That file can be read by a PostScript previewer (such as ghostscript) or sent to a printer (`lpr /tmp/chown.ps`). The next example creates HP LaserJet output (`-Tlj4`) and directs it to the default printer (`-l` option).

Creating a man page with Groff

Before HOW-TOs and info files, man pages were the foundation for information about UNIX (and UNIX-like) systems. Each command, file format, device, or other component either had its own man page or was grouped on a man page with similar components. To create your own man page requires that you learn a few macros (in particular, `man` macros). Figure 6-3 shows the source for a fictitious man page for a command called `waycool`.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
.\"
.\" waycool.1   the *roff document processor source for the waycool command
.\"
.TH waycool 1 "May 12, 2002" GNU "Linux Programmer's Manual"
.SH NAME
waycool \- my cool command
.SH SYNTAX
\fBwaycool\fR [ \fB-abcv\fR ] [ \fI file ... \fR ]
.SH VERSION
This man page documents the GNU waycool version X.XX
.SH DESCRIPTION
\fBwaycool\fR is a way cool command.
.SP
This version of \fBwaycool\fR is better than the last one.
.SH OPTIONS
.IP -a
Run all options with it.
.IP -b
Run some options.
.IP -c
affect symbolic links instead of any referenced file
(available only on systems that can change the
ownership of a symlink)
.IP -v
Print the version number with the command.
.SH COMMENTS
If you don't like the command, don't tell me. It will just hurt my feelings.
.SH ENVIRONMENT VARIABLES
These environment variables are used by  \fBwaycool\fR:
.IP "DISPLAY"
This sets the X Display variable.
.IP "WAYCOOL"
This contains the location of the waycool database.
.SH FILES
/usr/local/waycool - Directory containing waycool stuff.
.SH AUTHOR
Chris Craft <chris@handsonhistory.com>
.SH "REPORTING BUGS"
Report bugs to <bug-fileutils@gnu.org>.
.SH COPYRIGHT
Copyright \(co 2001 Free Software Foundation, Inc.
.br
This is free software; see the source for copying conditions.  There is NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
.SH ACKNOWLEDGEMENTS
I'd like to thank all my friends.
```

Figure 6-3: Simple markup is required to create man pages.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Most man pages are stored in subdirectories of /usr/share/man. Before you create a man page, refer to similar man pages to see the markup and the headings they include. In ma

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

A few other kinds of macros are used in the man page. The .IP macros format indented paragraphs for things such as options. The man page also contains some lower-level font requests; for example, \fB says to change the current font to bold, \fI changes the font to italic, and \fR changes it back to regular font. (This markup is better than asking for a particular font type because it just changes to bold, italic, or regular for the current font.) Figure 6-4 shows what the waycool man page looks like after it is formatted with `groff`:

waycool(1) Linux Programmer's Manual waycool(1)

NAME
waycool – my cool command

SYNTAX
waycool [-abcv] [*file* ...]

VERSION
This man page documents the GNU waycool version X.XX

DESCRIPTION
waycool is a way cool command. This version of waycool is better than the last one.

OPTIONS
-a Run all options with it.
-b Run some options.
-c affect symbolic links instead of any referenced file (available only on systems that can change the ownership of a symlink)
-v Print the version number with the command.

COMMENTS
If you don't like the command, don't tell me. It will just hurt my feelings.

ENVIRONMENT VARIABLES
These environment variables are used by waycool:
DISPLAY
This sets the X Display variable.
WAYCOOL
This contains the location of the waycool database.

FILES
/usr/local/waycool - Directory containing waycool stuff.

AUTHOR
Chris Craft <chris@hardcorehistory.com>

REPORTING BUGS
Report bugs to <bug-filesutils@gnu.org>.

COPYRIGHT
Copyright © 2001 Free Software Foundation, Inc.
This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

ACKNOWLEDGEMENTS
I'd like to thank all my friends.

Figure 6-4: Man page formatting adds headers and lays out the page of text.

$ **groff -man -Tps -l waycool.1**

Table 6-1 lists the macros that you can use on your man pages. These macros are described on the man(7) manual page (type **man 7 man** to view that page).

Table 6-1: Man Macros

Macro	Description
.B	Bold
.BI	Bold, then italics (alternate)
.BR	Bold, then roman (alternating)
.DT	Set default tabs
.HP	Begin a hanging indent
.I	Italics
.IB	Italics, then bold (alternating)
.IP	Begin hanging tag. For options. Long tags use .TP.
.IR	Italics, then roman (alternating)
.LP	Begin paragraph
.PD	Set distance between paragraphs
.PP	Begin paragraph
.RB	Roman, then bold (alternating)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 6-1: Man Macros

Macro	Description
.RE	End relative indent (after .RS)
.RI	Roman, then italics (alternating)
.RS	Begin relative indent (use .RE to end indent)
.SB	Small text, then bold (alternating)
.SM	Small text. Used to show words in all caps.
.SH	Section head
.SS	Subheading within a .SH heading.
.TH	Title heading. Used once at the beginning of the man page.
.TP	Begin a hanging tag. Begins text on next line, not same line as tag.

Creating a letter, memo, or white paper with Groff

Memorandum macros (which are used with the `-mm` option of Groff) were once popular among UNIX users for producing technical documents, letters, and memos. Although more modern word processors with a variety of WYSIWYG templates have made `mm` outdated, in a pinch `mm` can still be a quick way to create a typeset-style document in a text environment.

To format and print (to a PostScript printer) a document with `mm` macros, use the following:

$ **groff -mm -Tps -l letter.mm**

The following is a simple example of how to use `mm` macros to produce a letter:

```
.WA "Christopher T. Craft"
999 Anyway Way
Anytown, UT 84111 USA
.WE
.IA
John W. Doe
111 Notown Blvd.
Notown, UT 84111
.IE
.LO RN "Our telephone conversation"
.LO SA "Dear Mr. Doe:"
.LT
In reference to our telephone conversation on the 4th, I am calling to
confirm our upcoming appointment on the 18th. I look forward to
discussing the merger. I believe we have a win-win situation here.
.FC "Yours Truly,"
.SG
```

The output of the letter, if you use the `groff` command line mentioned in the paragraph preceding the code example, is shown in Figure 6-5.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

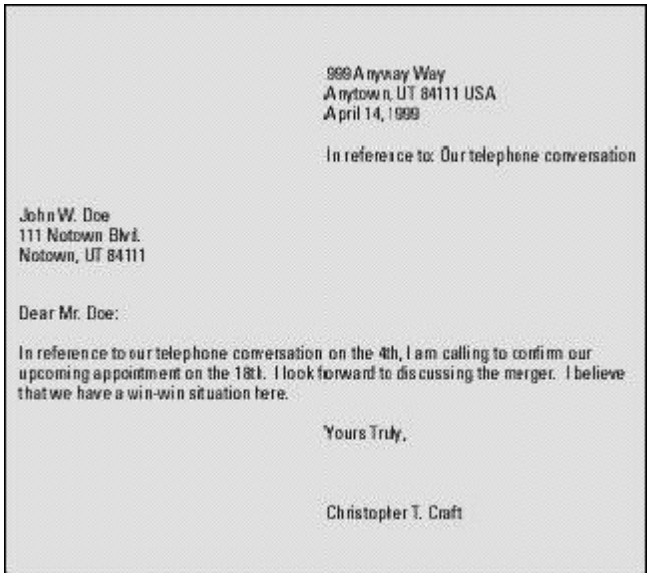

Figure 6-5: Create a simple letter using mm macros.

The `mm` macros were often used to produce technical memos. The following is an example of a sign-off sheet that might go at the front of a larger technical memo.

```
.TL
Merger Technical Specifications
.AF "ABC Corporation"
.AU "Christopher Craft"
.AT "President"
.AS
This memo details the specifications for the planned merger.
.AE
.MT "Merger Description and Marching Orders"
As a result of our talks with XYZ corporation, we plan to go
forward with the merger. This document contains the following:
.BL
.LI
Schedule and time tables.
.LI
Financial statements.
.LI
Asset allocations.
.LE
.SP
Please add any corrections you have, then sign the approval line
indicated at the bottom of this sheet.
.FC
.SG
.AV "John W. Doe, XYZ Corporation President"
.AV "Sylvia Q. Public, XYZ Corporation CFO"
.NS
Everyone in the corporation.
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

.NE

Figure 6-6 shows the output of this memo.

subject: Merger Technical Specification

date: April 11, 1999
from: Christopher Craft

ABSTRACT

This memo details the specifications for the planned merger.

Merger Description and Marching Orders

As a result of our talks with XYZ Corporation, we plan to go forward with the merger. This document contains the following:

- Schedule and time tables.
- Financial statements.
- Asset allocations.

Please add any corrections you have, then sign the approval line indicated at the bottom of this sheet.

Yours Very Truly,

Christopher Craft

APPROVED:

_____ _____
John W. Doe, XYZ Corporation President Date

APPROVED:

_____ _____
Sylvia Q. Public, XYZ Corporation CFO Date

copy to:
Everyone in the corporation.

Figure 6-6: Add headings and approval lines automatically to memos.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note For a complete listing of m m macros, see the groff_mman page. More than 100 m m macros exist. Also, dozens of defined strings let you set and recall information (such as

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Adding equations, tables, and pictures

To interpret special macros for equations, tables, and line drawings, you can run separate commands (`eqn`, `tbl`, and `pic` commands) on the file before you run the `groff` command. Alternatively, you can add options to the `groff` command line to have the file preprocessed automatically by any of the commands (`-e` for `eqn`, `-t` for `tbl`, and `-p` for `pic`).

Here are some examples of EQN, TBL, and PIC markup included in a Groff document. The first example shows an equation that can be processed by `eqn` for a Groff document:

```
.EQ
a ~ mark = ~ 30
.EN
.sp
.EQ
a sup 2 ~ + ~ b sup 2~lineup = ~ 1000
.EN
.sp
.EQ
x sup 3 ~ + ~ y sup 3 ~ + ~ z sup 3~lineup = ~ 1400
.EN
```

If this appeared in a memo called `memoeqn.mm`, the memo would be preprocessed by `eqn` and then sent to the printer using the following command:

```
$ groff -Tps -l -mm -e memoeqn.mm
```

All data between the .EQ and .EN macros are interpreted as equations. The resulting output from the equation would appear as shown in Figure 6-7.

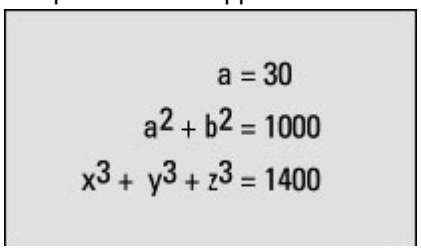

Figure 6-7: Produce equations in documents with the use of the eqn command's .EQ and .EN macros.

To create a table in a Groff document, use the .TS and .TE macros of the `tbl` preprocessor. The following is an example of the markup used to produce a simple table.

```
.TS
center, box, tab(:);
c s s
c | c | c
l | l | l.
Mergers and Acquisitions Team
=
Employee:Title:Location
=_
Jones, James:Marketing Manager:New York Office
Smith, Charles:Sales Manager:Los Angeles Office
Taylor, Sarah:R&D Manager:New York Office
Walters, Mark:Information Systems Manager:Salt Lake City Office
Zur, Mike:Distribution Manager:Portland Office
.TE
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

After the .TS macro starts the table, the next line indicates that the table should be centered on the page (`center`) and surrounded by a line box and that a colon will be used to separate the data into cells (`tab(:)`). The next line shows that the heading should be centered in the box (`c`) and should span across the next two cells (`s s`). The line after that indicates that the heading of each cell should be centered (`c | c | c`) and that the data cells that follow should be left justified (`l | l | l`).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution There must be a period at the end of the table definition line. In this case, it is after the l | l | l. line. If the period is not there, tbl will try to interpret the text as part of the table definition. In this

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The rest of the information in the table is the data. Note that the tab separators are colon characters (`:`). When the table is done, you end it with a .TE macro. If the table were in a memo called `memotbl.mm`, `tbl` could preprocess the memo and then send it to the printer using the following command:

$ **groff -Tps -l -mm -t memotbl.mm**

Data between .TS and .TE macros are interpreted as tables. Figure 6-8 displays this example.

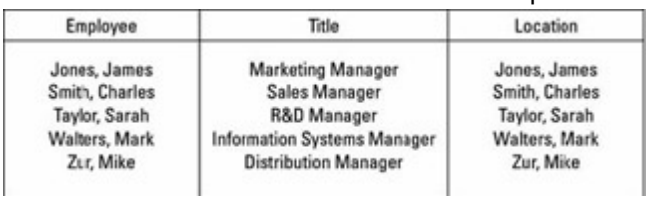

Employee	Title	Location
Jones, James	Marketing Manager	Jones, James
Smith, Charles	Sales Manager	Smith, Charles
Taylor, Sarah	R&D Manager	Taylor, Sarah
Walters, Mark	Information Systems Manager	Walters, Mark
Zur, Mike	Distribution Manager	Zur, Mike

Figure 6-8: Set how text is justified and put in columns with the use of the tbl command's .TS and .TE macros.

The PIC macros (.PS and .PE) let you to create simple diagrams and flow charts to use in Groff. PIC is really only qualified to create simple boxes, circles, ellipses, lines, arcs, splines, and some text. The following is some PIC code that could be in a Groff document:

```
.PS
box invis "Start" "Here"; arrow
box "Step 1"; arrow
circle "Step 2"; arrow
ellipse "Step 3"; arrow
box "Step 4"; arrow
box invis "End"
.PE
```

After the .PS, the first line indicates an invisible box (`invis`) that contains the words Start Here, followed by an arrow. That arrow connects to the next box containing the words Step 1. The next elements (connected by arrows) are a circle (Step 2), an ellipse (Step 3), another box (Step 4), and another invisible box (End). The .PE indicates the end of the pic drawing.

If these lines appeared in a document called memopic.mm, you could preprocess the PIC code and print the file using the following command:

$ **groff -Tps -l -mm -p memopic.mm**

Figure 6-9 shows an example of this drawing.

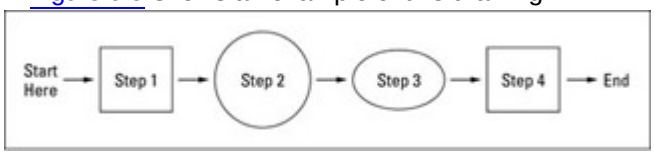

Figure 6-9: Create simple flow diagrams with the pic command's .PS and .PE macros.

Text processing with TeX/LaTeX

TeX is a collection of commands used primarily to produce scientific and mathematical typeset documents. The most common way to use TeX is by calling a macro package. The most popular macro package for Tex is LaTeX, which takes a higher-level approach to formatting TeX documents. TeX and LaTeX tools are contained in the tetex-latex package.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The tetex-* packages needed to use the TeX examples shown in this book are found on CD #2 that accompanies this book. The tetex-doc package is on CD #3.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

TeX interprets the LaTeX macros from the latex format file (`latex.fmt`). By default, the `latex.fmt` and `plain.fmt` format files are the only ones that are automatically built when the TeX package is installed. Other macro files that you can use with TeX include:

amstex — Mathematical publications, including The American Mathematical Society uses this as their official typesetting system.

eplain — Includes macros for indexing and table of contents.

texinfo — Macros used by the Free Software Foundation to produce software manuals. Text output from these macros can be used with the Red Hat Linux `info` command.

You can create a TeX/LaTeX file using any text editor. After the text and macros are created, you can run the `tex` command (or one of several other related utilities) to format the file. The input file is in the form `filename.tex`. The output is generally three different files:

`filename.dvi` — This is the device-independent output file that can be translated for use by several different types of output devices (such as PostScript).

`filename.log` — This is a log file that contains diagnostic messages.

`filename.aux` — This is an auxiliary file used by LaTeX.

The .dvi file produced can be formatted for a particular device. For example, you could use the `dvips` command to output the resulting .dvi file to your PostScript printer (`dvips filename.dvi`). Or you could use the `xdvi` command to preview the dvi file in X.

Creating and formatting a LaTeX document

Because LaTeX is the most common way of using TeX, this section describes how to create and format a LaTeX document. A LaTeX macro (often referred to as a command) appears in a document in one of the two following forms:

***string*{*option*}[*required*]** — First there is a backslash (`\`), which is followed by a string of characters. (Replace *`string`* with the name of the command.) Optional arguments are contained in braces (`{}`), and required arguments are in brackets (`[]`).

***?*{*option*}[*required*]** — First there is a backslash (`\`), which is followed by a single character that is not a letter. (Replace *?* with the command character.) Optional arguments are contained in braces (`{}`), and required arguments are in brackets (`[]`).

Each command defines some action to be taken. The action can control page layout, the font used, spacing, paragraph layout, or a variety of other actions on the document. The minimum amount of formatting that a LaTeX document can contain is the following:

```
\documentclass{name}
\begin{document}
  TEXT GOES HERE!
\end{document}
```

You should replace `{name}` with the name of the class of document you are creating. Valid document classes include article, book, letter, report, and slides. The text for the file, along with your formatting commands, goes between the `begin` and `end` document commands.

The best way to get started with LaTeX is to use the LyX editor. LyX provides a GUI for creating LaTeX documents. It also contains a variety of templates you can use instead of just creating a document from scratch. Figure 6-10 shows an example of the LyX editor.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

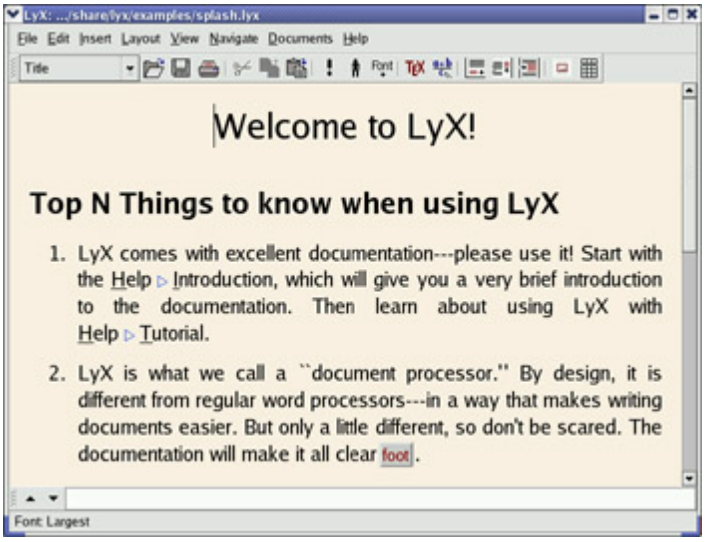

Figure 6-10: Create LaTeX documents graphically with the LyX editor.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The LyX editor doesn't come with Red Hat Linux. Find an RPM package for LyX from the LyX site at ftp://ftp.lyx.org/pub/lyx/bin. Look for a Red Hat RPM.

If you want to edit LaTeX in a regular text editor, you need to be familiar with the LaTeX commands. For a complete listing of the LaTeX commands, type `info latex` and then go to the section "Commands within a LaTeX document."

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the LyX LaTeX Editor

Start the LyX LaTeX editor with the `lyx` command. LyX comes with a lot of supporting documentation. Click Help to select a Tutorial, User's Guide, or other information.
To start your first document, I recommend that you select one of the templates provided with LyX. Templates are located in `/usr/share/lyx/templates`. To open a template, click File → New from Template. A list of available templates appears. You can use them to create letters, slides, and articles, for example.

Besides offering standard editing functions, such as cut, copy, and paste, you can perform a variety of markup functions from the Layout menu. As for mathematical functions, the Math menu enables you to insert fractions, square root, exponent, sum, and integral functions into your document. When you are done, you can:

Print the file to a PostScript printer or output a PostScript (`.ps`) file. (Click File → Print, select the printing method, and then click OK.)
Export the file to LaTeX, DVI, PostScript, or ASCII Text. (Click File → Export and choose from the list of file formats.)

LyX calls itself a WYSIWYM editor — What You Say Is What You Mean. As a result, what you see on the screen as you edit is not exactly what the printed document will look like. For example, no extra white space will appear between lines by pressing Enter multiple times.

Because LyX supports style files, it enables you to create documents that meet several different standards. For example, LyX supports typesetting for the American Mathematics Society (AMS) journals using the article text class. Other text classes supported include:

article — One-sided paper with no chapters.
report — Two-sided report, tending to be longer than an article.
book — Same as report, with additional front and back matter.
slides — For producing transparencies.
letter — Includes special environments for addresses, signatures, and other elements.

Printing LaTeX files

Whether you create your own LaTeX file, export one from the LyX LaTeX editor, or download one from the Internet, several utilities are available to format, print, or display the output. Here are some of your choices:

To format a LaTeX file (`filename.tex`), run the following command:
$ **latex *filename.*tex**
To print a .dvi file (`filename.dvi`), send it to your default PostScript printer, and type the following:
$ **dvips *filename.*dvi**
To display a .dvi file in an X window, type the following:
$ **xdvi *filename*.dvi**
To print a .dvi file to a PCL printer, such as an HP LaserJet, type the following:
$ **dvicopy *filename*.dvi**
$ **dvilj filename.dvi**
The `dvilj` command doesn't support virtual fonts directly. The `dvicopy` command converts the fonts so that the PCL printer can handle them.

Converting documents

Documents can come to you in many different formats. Search just some of the Linux FTP sites on the Internet and you will find files in PostScript, DVI, man, PDF, HTML, and TeX. There are also a variety of graphics formats. Red Hat Linux comes with lots of utilities to convert documents and graphics from one format to another. The following is a list of document and graphics conversion utilities:

dos2unix — Converts a DOS text file to a UNIX (Linux) text file.
fax2ps — Converts TIFF facsimile image files to a compressed PostScript format. The PostScript output is optimized to send to a printer on a low-speed line. This format is less efficient for images with a lot of black or continuous tones. (In those cases, tiff2ps might be more effective.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

fax2tiff — Converts fax data (Group 3 or Group 4) to a TIFF format. The output is either low-resolution or medium-resolution TIFF format.
g32pbm — Converts a Group 3 fax file (either digifax or raw) to a portable bitmap.
gif2tiff — Converts a GIF (87) file to a TIFF format.
man2html — Converts a man page to an HTML format.
pal2rgb — Converts a TIFF image (palette color) to a full-color RGB image.
pbm2g3 — Converts a portable bitmap image to a fax file (Group 3).
pdf2dsc — Converts a PDF file to a PostScript document dsc file. The PostScript file conforms to Adobe Document Structuring Conventions (DSC). The output enables PostScript readers (such as Ghostview) to read the PDF file a page at a time.
pdf2ps — Converts a PDF file to a PostScript file (level 2).
pfb2pfa — Converts Type 1 PostScript font (binary MS-DOS) to ASCII-readable.
pk2bm — Converts a TeX pkfont font file to a bitmap (ASCII file).
ppm2tiff — Converts a PPM image file to a TIFF format.
ps2ascii — Converts PostScript or PDF files to ASCII text.
ps2epsi — Converts a PostScript file to Encapsulated PostScript (EPSI). Some word processing and graphic programs can read EPSI. Output is often low quality.
ps2pdf — Converts PostScript file to Portable Document Format (PDF).
ps2pk — Converts a Type 1 PostScript font to a TeX pkfont.
pstotext — Converts a PostScript file to ASCII text. pstotext is similar to ps2ascii, but handles font encoding and kerning better than ps2ascii. pstotext doesn't convert PDFs.
ras2tiff — Converts a Sun raster file to a TIFF format.
texi2html — Converts a Texinfo file to HTML.
tiff2bw — Converts an RGB or Palette color TIFF image to a grayscale TIFF image.
tiff2ps — Converts a TIFF image to PostScript.
unix2dos — Converts a UNIX (Linux) text file to a DOS text file.

Besides these tools many graphical applications, such as The GIMP, let you save images into several different formats (BMP, JPEG, PNG, TIFF, and so on), using the Save As feature.

Creating DocBook documents
Documentation projects often need to produce documents that are output in a variety of formats. For example, the same text that describes how to use a software program may need to be output as a printed manual, an HTML page, and a PostScript file. The standards that have been embraced most recently by the Linux community for creating what are referred to as *structured documents* are SGML, XML, and DocBook.

Understanding SGML and XML

Standard Generalized Markup Language (SGML) was created to provide a standard way of marking text so that it could be output later in a variety of formats. Because SGML markup is done with text tags, you can create SGML documents using any plain-text editor. Documents consist of the text of your document and tags that identify each type of information in the text.

Unlike markup languages such as Groff and HTML, SGML markup is not intended to enforce a particular look when you are creating the document. So, for example, instead of marking a piece of text as being bold or italic, you would identify it as an address, paragraph, or a name. Later, a style sheet would be applied to the document to take the tagged text and assign a look and presentation.

Because SGML consists of many tags, to simplify producing documents based on SGML other projects have cropped up to better focus the ways in which SGML is used. In particular, the Extensible Markup Language (XML) was created to offer a manageable subset of SGML that would be specifically tailored to work well with Web-based publishing.

So far in describing SGML and XML, I have only referred to the frameworks that are used to produce structured documents. Specific documentation projects need to create and, to some extent, enforce specific markup definitions for the type of documents they need to produce. These definitions are referred to as Data Type Definitions (DTDs). For documentation of Linux itself and other open source projects, DocBook has become the DTD of choice.

Understanding DocBook

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

DocBook is a DTD that is well suited for producing computer software documents in a variety of formats. It was originally created by the OASIS Consortium (www.oasis-open.org) and is now supported by many different commercial and open-source tools.

DocBook's focus is on marking content, instead of indicating a particular look (that is, font type, size, position, and so on.). It includes markup that lets you automate the process of creating indices, figure lists and tables of contents, to name a few. Tools in Red Hat Linux let you output DocBook documents into HTML, PDF, DVI, PostScript, RTF, and other formats.

DocBook is important to the Linux community because many open-source projects are using DocBook to produce documentation. For example, the following is a list of organizations, and related Web sites, that use DocBook to create the documents that describe their software:

Linux Documentation Project (www.tldp.org/LDP/LDP-Author-Guide)
GNOME Documentation (
developer.gnome.org/projects/gdp/handbook/gdp-handbook)
KDE Documentation Project (www.kde.org/documentation)
FreeBSD Documentation Project (www.freebsd.org/docproj)

If you want to contribute to any of the above documentation projects, refer to the Web sites for each organization. In all cases, they publish writers' guides or style guides that describe the DocBook tags that they support for their writing efforts.

Creating DocBook documents

You can create the documents in any text editor, using tags that are similar in appearance to HTML tags (with beginning and end tags appearing between less-than and greater-than signs). There are also word processing programs that allow you to create DocBook markup.

The following procedure contains an example of a simple DocBook document produced with a plain-text editor and output into HTML using tools that come with Red Hat Linux.

1. Create a directory in your home directory to work in and go to that directory. For example, you could type the following from a Terminal window:
2. $ **mkdir $HOME/doctest**
3. $ **cd $HOME/doctest**
4. Open a text editor to hold your DocBook document. For example, you could type:
5. $ **kedit cardoc.sgml**
6. Enter the tags and text that you want to appear in your document. Most DocBook documents are either <book> type (large, multi-chapter documents) or <article> type (single chapter documents). To try out a DocBook document, type the following:
7. <xml version="1.0">
8. <article>
9. <title>Choosing a new car</title>
10. <artheader>
11. <abstract>
12. In this article, you will learn how to price,
13. negotiate for, and purchase an automobile.
14. </abstract>
15. </artheader>
16. <section>
17. <title>Getting Started</title>
18. <para>
19. The first thing you will learn is how to figure out
20. what you can afford.
21. </para>

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
</section>
<section>
  <title>The Next Step</title>
  <para>
   After you know what you can afford, you can begin your
   search.
  </para>
 </section>
</article>
```

There are a few things you should notice about this document. The entire document is wrapped in article tags (<article> </article>). The article title is in title tags (<title> </title>). The section tags (<section> </section>) indicate sections of text that have a title and paragraph each. These sections can later be treated separately in the TOC.

30. Save the file and exit from the text editor.
31. Next, you can try translating the document you just created into several different formats. For example, to create HTML output you could type the following:
32. $ **db2html cardoc.sgml**
    The result is a new directory called `cardoc`. The result from `db2html` in the `cardoc` directory was: stylesheet-images directory, `t2.html` file, and `x12.html` file.
33. To view the HTML file just created, I typed the following:
34. $ **epiphany $HOME/doctest/cardoc/t2.html**
    Figure 6-11 shows an example of the output created from the `db2html` command. The screen on the left shows the first page. Click on the Next link at the top of the page. The second page that you see is shown on the right. During conversion to HTML, the `db2html` command adds Next/Previous buttons to each page. It also puts the title of each section in a Table of Contents on page one and in the browser's title bar.

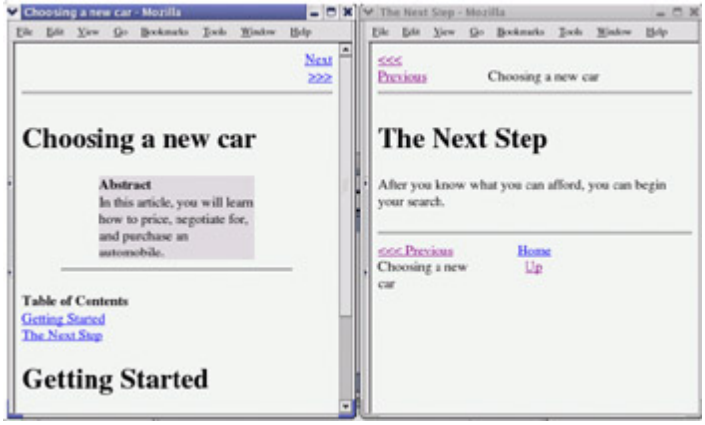

Figure 6-11: The DocBook file is output in HTML with the db2html command.

From this point, you can continue to add content and different types of tags. If you are writing documents for a particular project (such as the Linux projects mentioned earlier), you should get information on the particular tags and other style issues they require.

Converting DocBook documents

The previous example shows how to create a simple DocBook document and convert it to HTML output. The following Red Hat Linux utilities convert DocBook to other formats:

docbook2dvi — Converts a DocBook file to Device Independent file format.
docbook2html — Converts a DocBook file to HTML format.
docbook2man — Converts a DocBook file to man page format.
docbook2pdf — Converts a DocBook file to Portable Document Format (PDF).
docbook2ps — Converts a DocBook file to PostScript format.
docbook2rtf — Converts a DocBook file to Rich Text Format (RTF).
docbook2tex — Converts a DocBook file to TeX format.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

docbook2texi — Converts a DocBook file to GNU TeXinfo format.
docbook2txt — Converts a DocBook file to a bare text format.

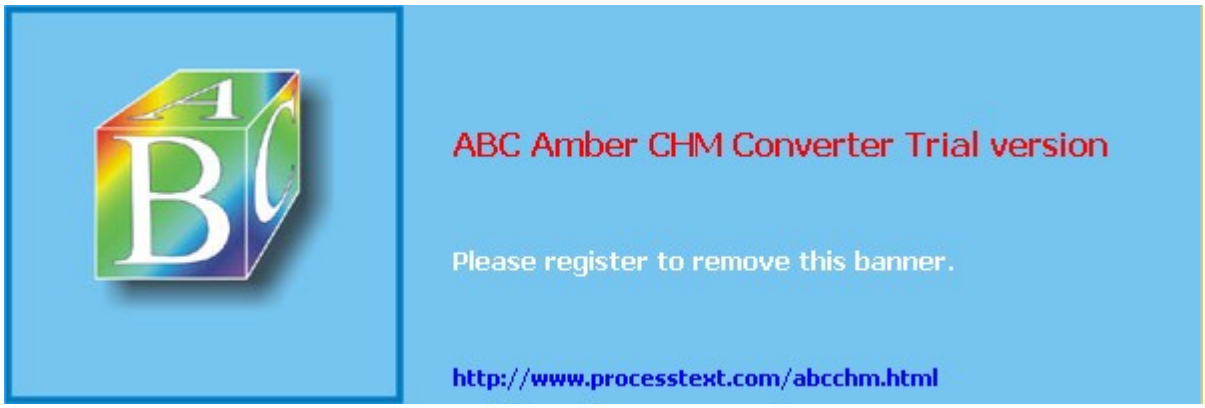

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Printing Documents with Red Hat Linux

Printing in Red Hat Linux is provided by the Common UNIX Printing System (CUPS) service. The LPRng service, which is no longer the default printing service used by Red Hat Linux, still can be obtained from Red Hat FTP sites. Both services, however, let you print using the same basic set of printing commands described in this section.

Cross-Reference See Chapter 10 for a description of how to use the Alternatives system to choose between LPRng and CUPS as the default print service. For information on configuring local and remote printers for Red Hat Linux, see Chapter 17.

As a nonadministrative user, you don't have a lot of control over how the printers are configured. You can, however, check which printers are available to print to, check the status of print queues (documents waiting to print), and remove any of your own queued print jobs.

Printing to the default printer

When your system administrator (or you) configured printers for your computer, one of those printers was defined as the default printer. If you are not sure which printer is your default, type `redhat-config-printer` and look for the printer with the check by it.

Most graphical word processors, such as StarOffice and OpenOffice, let you choose a printer from those available. Some of the less sophisticated Linux utilities that run from the command line, however, use only the default printer. For example, `dvips` (to print a PostScript file) and `groff -l` (to print a troff/nroff file) automatically send the output to the default printer.

As a regular user, you can override the default printer using the `PRINTER` environment variable. For example, if the default printer on your computer is `lp0` and you want to print regularly to `lp1`, change your default printer by setting the `PRINTER` variable as follows:

$ **export PRINTER=lp1**

To have this take effect all the time, you could add this line to one of your shell configuration files (such as `$HOME/.bashrc`, if you use the bash shell).

Printing from the shell

The `lpr` command is used to print files from the shell. You can use `lpr` to print whether the LPRng or CUPS print service is being used. If you have a file already formatted, use `lpr` to print it. For example, if you have a PostScript output file (file.ps) and you want to print it to your PostScript printer, use the following command line:

$ **lpr file.ps**

If you want to specify a particular printer (other than the default), add the `-P`*`printer`* option. For example, to print to the lp0 printer, you could type the following:

$ **lpr -Plp0 file.ps**

If you want to print more than one copy of a document, use the `-#`*`num`* option, where *`num`* is replaced by the number of copies you want. For example, to print five copies of a file, use:

$ **lpr -#5 file.ps**

The `lpr` command can also accept standard output for printing. For example, you could print the output of a `groff` command by piping that output to `lpr` as follows:

$ **groff -Tps -man /tmp/chown.1 | lpr -Plp0**

Tip The enscript command (in the enscript package) is another useful tool for printing plain-text files. It converts the files to PostScript and sends them to a printer or to a specified file.

Checking the print queues

To check the status of print jobs that have been queued, you can use the `lpq` command. By itself, `lpq` prints a listing of jobs that are in the queue for the default printer. For example:

$ **lpq**

hp is ready and printing

Rank	Owner	Job	Files	Total Size
active	root	3	hosts	1024 bytes

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
1st     root      7    (stdin)       625 bytes
2nd     root      8    memo1.ps      12273 bytes
3rd     chuck     9    bikes.ps      10880 bytes
```

The output from `lpq` shows the printer status and the files waiting to be printed. `Rank` lists the order in which they are in the queue. `Owner` is the user who queued the job. `Job` shows the job number. The `Files` column shows the name of the file or standard output (if the file was piped or directed to `lpr`). `Total Size` shows how large each file is in bytes.

You can add options to `lpq` to print different kinds of information. By adding `-P`*`printer`*, you can see the queue for any available printer. You can also add the job number (to see the status of a particular print job) or a user name (to see all queued jobs for a user).

Removing print jobs

If you have ever printed a large document by mistake, you understand the value of being able to remove a print job from the queue. Likewise, if a printer is going to be down for a while and everyone has already printed their jobs to another printer, it's sometimes nice to be able to clear all the print jobs when the printer comes back online.

Remove print jobs in Red Hat Linux using `lprm`. For example, to remove all jobs for the user named `bill` (assuming you are either `bill` or the root user), type the following:

$ **lprm bill**

The root user can remove all print jobs from the queue. To do this you add a dash (`-`) to the `lprm` command line as follows:

$ **lprm -**

You can also remove queued print jobs for a particular printer (`-P`*`printer`*) or for a particular job number by just adding the job number to the `lprm` command line.

Checking printer status

Sometimes nothing comes out of a printer and you have no idea why. The `lpc` command is a printer status command you can run that might give you a clue as to what is going on with your printer. The `lpc` command is intended for administrators, so it may not be in your default PATH. To start the `lpc` command, type the following:

```
# /usr/sbin/lpc
lpc> status
hp:
        printer is on device 'lpd' speed -1
        queing is enabled
        printing is enabled
        no entries
        daemon present
lpc>
```

When the command returns the `lpc>` prompt, type the word `status`. This example shows the status of printer `hp`. Here, queuing and printing are enabled. The printer shows no problems, no print jobs are waiting.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Displaying Documents with Ghostscript and Acrobat

Red Hat Linux publishing can be very paper-intensive if you send a Groff or LaTeX document to the printer each time you want to make a change to the document's content or formatting. To save paper and time spent running around, you can use some print preview programs to display a document on the screen as it will appear on the printed page. The following sections describe the `ghostscript` command for displaying PostScript files and the Adobe Acrobat reader for displaying Portable Document Format (PDF) files.

Using the ghostscript and gv commands

To display PostScript or PDF documents in Red Hat Linux, you can use the `ghostscript` command. The `ghostscript` command is a fairly crude interface, intended to let you step through documents and interpret them one line at a time. (If the `ghostscript` command is not installed on your system, you can get it by installing the ghostscript package from CD #1 that comes with this book.)
You can display any `.ps` or `.pdf` file you happen to have on your computer. For example, if the samba package is installed, you could type the following to display a PDF file (otherwise, you could find your own PDF file to try it):

$ **ghostscript /usr/share/doc/samba-*/docs/Samba-HOWTO-Collection.pdf**

>>showpage, press <return> to continue<<
At the prompt, press Enter (or Return) to go through the file one page at a time. When the file is done, you can type the name of another PostScript or PDF file and page through that file. When you are done, type **quit**.
The `gv` command (ghostview) is another, more friendly way of viewing PostScript files. (If the `gv` command is not installed on your system, you can get it by installing the gv package from CD #3 that comes with this book.)
To use `gv` to open a file called `rbash.ps`, you would type the following:

$ **gv /usr/share/doc/bash-doc-*/bashref.ps**

When the ghostview window opens, you can see the document. Left-click your mouse on the page and move it up and down to scroll the document. Use Page Up and Page Down keys to page through the document. You can click on a page number in the left column to jump to a particular page or click on the Print All button to print the entire document.

Using Adobe Acrobat Reader

The Portable Document Format (PDF) provides a way of storing documents as they would appear in print. With Adobe Acrobat Reader, you can view PDF files in a very friendly way. Adobe Acrobat makes it easy to move around within a PDF file. A PDF file may include hyperlinks, a table of contents, graphics, and a variety of type fonts.
A recent version of the Adobe Acrobat Reader (version 5.06) is available in RPM format from Guru Labs (`www.gurulabs.com/downloads.html`). This version can use many PDF features that aren't available in other PDF readers, including some new compression features in version 1.4. (While you are at it, you can install the Acrobat Plug-in RPM, to use the same reader to play PDF content when you browse the Web).

After you install Adobe Acrobat Reader, type the following command to start the program:

$ **acroread**
Click File → Open, and then select the name of a PDF file you have to display. Figure 6-12 shows an example of a PDF file viewed in Adobe Acrobat Reader.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

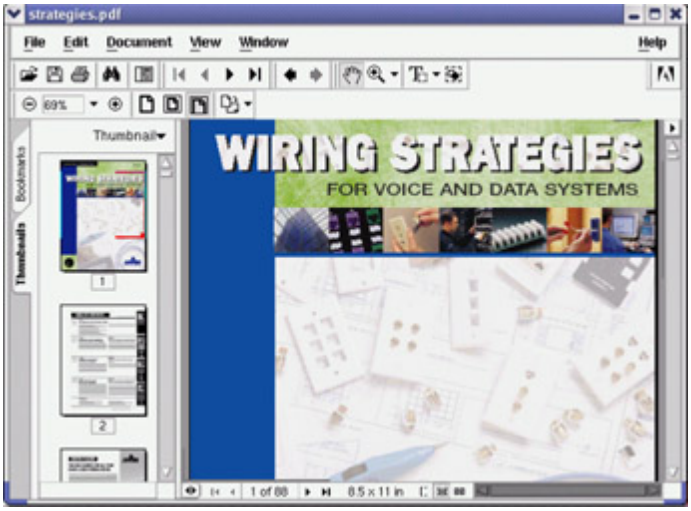

Figure 6-12: Display PDF files in the Adobe Acrobat Reader.

Acrobat Reader has a lot of nice features. For example, you can display a list of bookmarks alongside the document and click on a bookmark to take you to a particular page. You can also display thumbnails of the pages to quickly scroll through and select a page.

Using the menu bar or buttons, you can page through the PDF document, zoom in and out, go to the beginning or end of the document, and display different views of the document (as well as display bookmarks and page thumbnails). To print a copy, click File → Print.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version

Please register to remove this banner.

http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Working with Graphics

Tools for creating and manipulating graphics are becoming both more plentiful and more powerful in Red Hat Linux. Leading the list is the GNU Image Manipulation Program (GIMP). GIMP lets you compose and author images as well as retouch photographs. Other tools that come with Red Hat Linux for creating graphics include `ksnapshot` (a program for taking screen captures) and `kpaint` (for working with bitmap images).

Cross-Reference See Chapter 8 for descriptions of other multimedia applications, such as the gphoto window for working with images from digital cameras.

Manipulating images with GIMP

The GIMP is a free software program that comes with Red Hat Linux for manipulating photographs and graphical images. To create images with GIMP, you can either import a drawing, photograph, or 3D-image, or you can create one from scratch. You can start GIMP from the system menu by clicking Graphics → The GIMP or by typing **`gimp&`** from a Terminal window.

Figure 6-13 shows an example of GIMP.

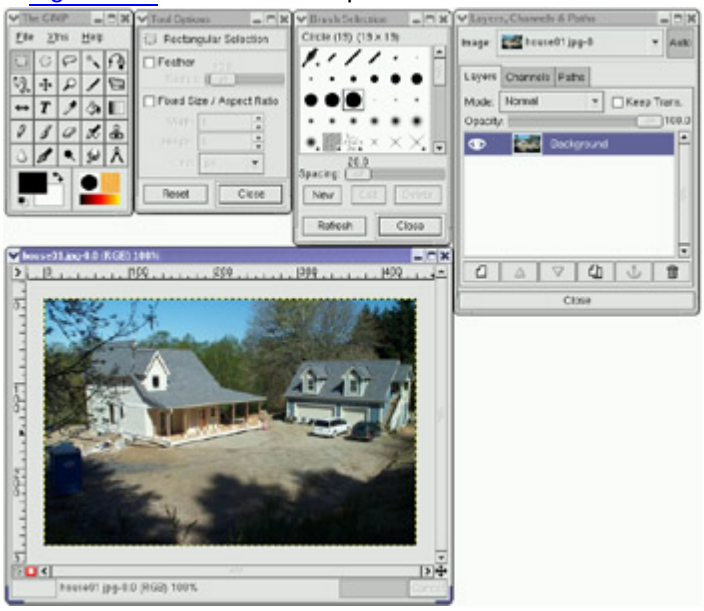

Figure 6-13: GIMP is a powerful tool for graphic manipulation.

Note If GIMP is not on your system or is not installed properly, you can install it from CD # 1 that comes with this book. Alternatively, you can obtain the latest copy of GIMP from www.gimp.org.

In many ways, GIMP is similar to Adobe Photoshop. Some people feel that GIMP's scripting features are comparable to, or even better than, Actions in Adobe Photoshop. One capability that GIMP lacks, however, is support for CMYK separations.

Taking screen captures

If you want to show examples of the work you do on Red Hat Linux, you can use the Screen Capture program to capture screen images. (The `ksnapshot` command is part of the kdegraphics package contained on CD #2.)

To open Screen Capture, from the Red Hat menu click Graphics → Screen Capture Program (or type `ksnapshot`). Figure 6-14 shows an example of the Screen Capture program.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

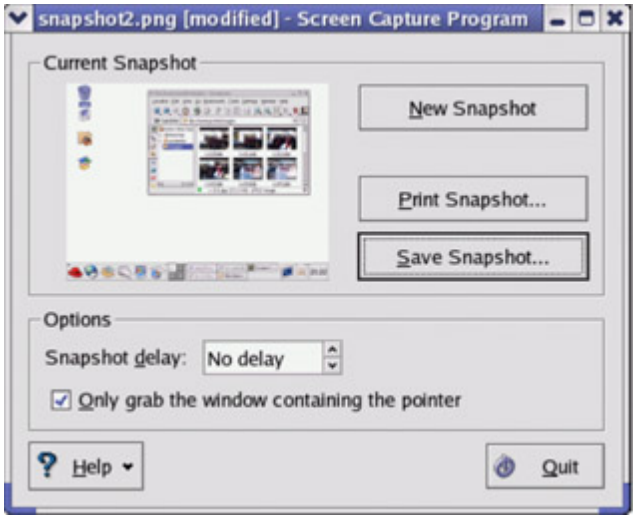

Figure 6-14: Grab a picture of your desktop or selected window with Screen Capture.

When Screen Capture first opens, it takes a snapshot of the full desktop. Buttons on the window let you:

Print Snapshot — Have the snapshot sent to your printer.

Save Snapshot — Save the snapshot to a file in X bitmap, windows icon, PNG, portable pixmap, JPEG, X pixmap, Encapsulated PostScript, or Windows BMP formats.

New Snapshot — Take a new snapshot. With "Only grab the window…" selected, you can click a window to capture that window alone.

Modifying images with KPaint

Using the KPaint window, a utility that comes with KDE (in the kdegraphics package on CD #2), you can work with and convert images in several formats. Figure 6-15 shows an example of KPaint.

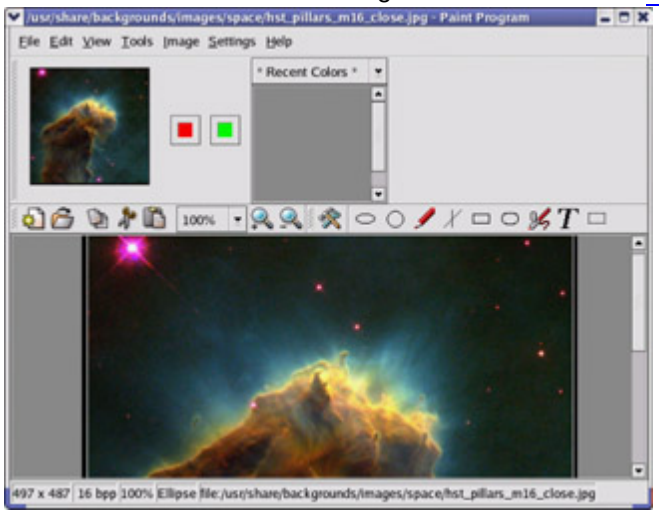

Figure 6-15: Edit bitmap images with KPaint.

Start KPaint from either the desktop (from the Red Hat menu, click Graphics→ More Graphics Applications → Paint Program) or from a Terminal window (`/usr/bin/kpaint&`). Start with either a blank canvas or by opening an image in one of the supported formats (File → Open, browse for a file, and then select OK). Look in the `/usr/share/backgrounds` directory for graphics to try.

The painting tools let you draw ovals, boxes, lines, and other shapes. You can save the file to several different formats, including MS Windows pixmaps, Encapsulated PostScript image, MS Windows icons, JPEG, PNG, PNM, TIFF, X bitmap, or X Window pixmap.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using Scanners Driven by SANE

Software for using a scanner with Linux is being driven by an effort called Scanner Access Now Easy (SANE). This effort hopes to standardize how device drivers for equipment such as scanners, digital still cameras, and digital video cameras are created, as well as help simplify the interfaces for applications that use those devices.
SANE is now included with the Red Hat Linux distribution. The sane-backend and sane-frontend packages are on CD #1, while the xsane and xsane-gimp packages are on CD #2. You can get the latest SANE driver packages from `www.mostang.com/sane`.

Someone wanting to use Linux as a publishing platform is generally interested in two issues about scanners: which scanners are supported and which applications are available to use the scanners. In general, more SCSI scanners are supported than parallel scanners.
Because of the ongoing development effort, new scanners are being supported all the time. You can find a current list of supported scanners at `www.mostang.com/sane/sane-backends.html`. As for applications, these are currently available with Red Hat Linux:

xsane — This is an X-based graphical front end for SANE scanners; `xsane` can work as a GIMP plug-in or as a separate application (from the Red Hat menu, select Graphics → Scanning). It supports 8-bit output in JPG, TIFF, PNG, PostScript, and PNM formats. There is experimental 16-bit support for PNM (ASCII), PNG, and raw formats.

scanimage — This is a command-line interface for obtaining scanned images. It supports the same formats as `xscanimage`. The command acquires the scanned image, and then directs the data to standard output (so you can send it to a file or pipe it to another program).

Because of the architecture of SANE scanner drivers, it is possible to separate scanner drivers from scanner applications. This makes it possible to share scanners across a network.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

In recent times, modern GUI-based publishing tools have augmented the text-based publishing tools that have always been available with Red Hat Linux (and other UNIX systems). Powerful open-source publishing tools such as OpenOffice are becoming competitive with commercial office suites. Traditional publishing tools such as Groff (which implements traditional troff/nroff text processing) and LaTeX (a TeX macro interface that is particularly suited for scientific and mathematical publishing) are still available with Red Hat Linux.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 7: Playing Games with Red Hat Linux

Overview

The advancement of computer games has mirrored the improvements in computers themselves. In the 1970s, the first games for UNIX systems were visually simple and could run on slow, character-based terminal connections. Today, games that combine graphics, animation, and sound have helped drive improvements in computer technology in general.

Availability of gaming software useable with Red Hat Linux is similar to that of Linux publishing software. A lot of the old software is still around (and is free), while newer software is available in demo form but costs some money to get a full version. Some experts predict that gaming will be the software category that brings Linux into homes. Although the number of popular game applications is limited at the moment, like everything else in Linux, more are becoming available each day.

This chapter addresses the current state of gaming in Linux, including the basics on getting your gaming environment going, and hardware considerations for gaming. It describes the free games (mostly fairly simple X Window games) that come with Red Hat Linux or that can be easily downloaded. For running games that were created for other platforms, this chapter describes game emulators such as WineX.

This chapter also discusses some popular commercial games that have demo versions available for Linux. If you like the demos, you can purchase these games, which run natively in Linux.

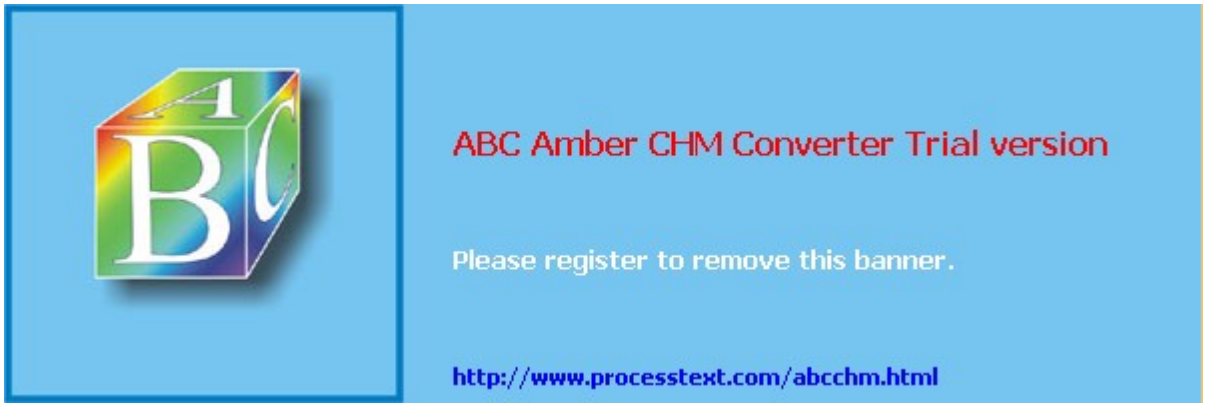

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Basic Linux Gaming Information

There isn't much you need to know to run most of the X Window-based games that come with Red Hat Linux. The following sections describe basic information about Linux gaming.

Where to get information on Linux gaming

To find news on the latest games available for Linux, as well as links to download sites, go to some of the several Web sites available. Here are a few to get you started:

TransGaming Technologies (www.transgaming.com) — This company's mission is to bring games from other platforms to Linux.

The Linux Game Tome (http://happypenguin.org) — This site features a database of descriptions and reviews of tons of games that run in Linux. You can do keyword searches for games listed at this site. The site also includes links to where you can get the different games, as well as links to other gaming sites.

Linuxgames.com (http://linuxgames.com) — This site can give you some very good insight into the state of Linux gaming. There are links to HOW-TOs and Frequently Asked Questions (FAQs), as well as forums for discussing Linux games. There are also links to Web sites that have information on a particular game you are interested in.

id Software (www.idsoftware.com) — Go to the id Software site for information on Linux demo versions for Quake, Unreal Tournament, and Return to Castle Wolfenstein.

Loki Entertainment Software (www.lokigames.com) — Loki provided ports of best-selling games to Linux, but went out of business in 2001. Its products included Linux versions of Civilization: Call to Power, Myth II: Soulblighter, SimCity 3000, Railroad Tycoon II, and Quake III Arena. The Loki Demo Launcher is still available to see demo versions of these games, and some boxed sets are available for very little money.

Tux Games (www.tuxgames.com) — If you are ready to purchase a game, the Tux Games Web site is dedicated to the sale of Linux games. Besides offering Linux gaming news and products, the site lists its top-selling games and includes notices of games that are soon to be released.

Linux Gamers' FAQ (http://icculus.org/lgfaq) — This FAQ contains a wealth of information about free and commercial Linux games. It lists gaming companies that have ported their games to Linux, tells where to get Linux games, and answers queries related to common Linux gaming problems.

If the idea of developing your own games interests you, try the Linux Game Development Center (http://lgdc.sunsite.dk).

Getting started with Linux gaming

How you get started with Linux gaming depends on how serious you are about it. If all you want to do is play a few games to pass the time, you can find plenty of diverting X Window games that come with Linux. If you want to play more powerful commercial games, you can choose from:

Games for Microsoft Windows (WineX) — Many of the most popular commercial games created to run on Microsoft operating systems will run in Linux using WineX. To get RPM versions of WineX, you must sign up for a WineX subscription at Transgaming.com.

Games for Linux (id Software and others) — Certain popular games have Linux versions available. Most notably, id Software offers its Unreal Tournament 2003, DOOM, and Return to Castle Wolfenstein in Linux versions.

Games are still available from the now defunct company Loki Software, Inc. I just purchased Myth II: Soulblighter and Heretic II for Linux over the Internet for a few dollars.

Choosing a video card for gaming

Because high-end games place extraordinary demands on your video hardware, choosing a good video card and configuring it properly is one of the keys to ensuring a good gaming experience. Although basic video card configuration is covered in Chapter 3, for advanced gaming you may need to go beyond

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

what a low-end card can do for you.
One feature that many games may require of your video card is Direct Rendering Infrastructure (DRI). Whether you are running the games using WineX or natively in Linux, to play demanding games in Linux you need a card that supports DRI to do hardware acceleration. Here is a list of video cards that support DRI from the DRI project site (http://dri.sourceforge.net/dri_status.phtml):

3dfx — Although 3dfx Interactive, Inc. is no longer in business, you can still find 3dfx cards that support DRI. In particular, the Voodoo (3, 4, and 5) and Banshee chip sets have drivers that support DRI. Voodoo 5 cards support 16 and 24 bpp. Scan Line Interleaving (SLI), where two or more 3D processors work in parallel (to result in higher frame rates), is not supported for 3dfx cards.

3Dlabs — Graphics cards containing the MX/Gamma chipset from 3Dlabs have drivers available that support DRI in Linux.

ATI Technologies — Chipsets from ATI Technologies that support DRI include the Mach64 (Rage Pro), Radeon 7X00 (R100), Radeon 2 / 8500 (R200), and Rage 128 (Standard, Pro, Mobility). Cards based on these chip sets include All-in-Wonder 128, Rage Fury, Rage Magnum, Xpert 99, Xpert 128, and Xpert 2000.

Intel — Supported video chipsets from Intel include the i810 (e, e2, and -dc100), i815 and i815e.

Matrox — The Matrox chipsets that have drivers that support DRI include the G200, G400, and G450. Cards that use these chips include the Millennium G450, Millennium G400, Millennium G200, and Mystique G200.

NVIDIA — Cards from NVIDIA are not supported by DRI because NVIDIA has not released hardware specifications to DRI developers. However, NVIDIA cards work for most Linux games. To get NVIDIA drivers, which are produced by NVIDIA but are not open source drivers, you must download them from the NVIDIA Web site. From the NVIDIA Web site (www.nvidia.com), click the download button and follow the instructions for downloading and installing the correct drivers for your card on Red Hat Linux systems. RPM packages are available.

To find out whether DRI is working on your current video card, type the following:

```
$ glxinfo | grep rendering
```

```
direct rendering: Yes
```

This example shows that direct rendering is supported. If it were not supported, the output would say `No` instead of `Yes`.

Team LiB

PREVIOUS NEXT

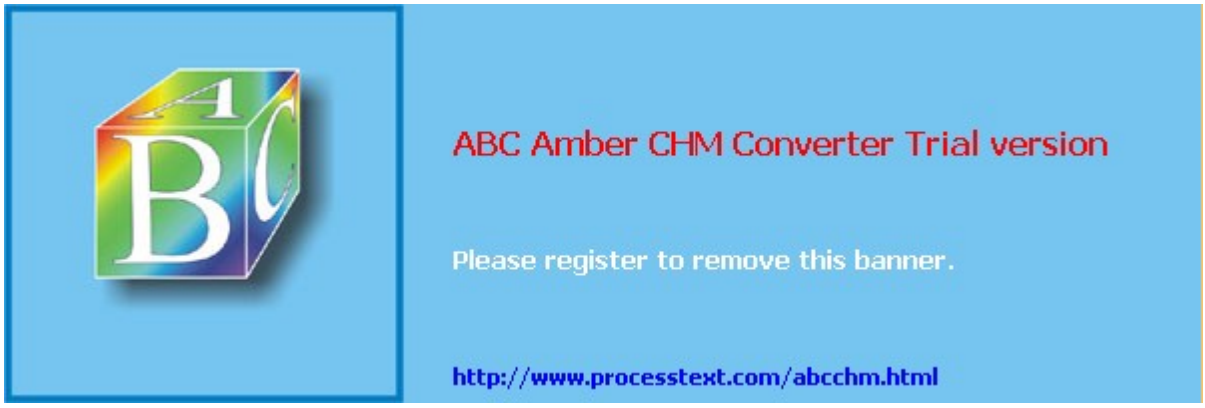

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

X Window Games

The X Window System created a great opportunity for games in Red Hat Linux and other UNIX systems to become graphical-based rather than character-based. So, instead of having little character symbols representing robots and arrows, the games could actually show pictures of little robots and arrows.
A lot of diverting games come with Red Hat Fedora and run in X. Unless otherwise noted, all of the X games described in this section are free. Also, the GNOME and KDE environments that come on the CDs (described in Chapter 3) each have a set of games associated with it.

GNOME games

The GNOME games consist of some old card games and a bunch of games that look suspiciously like games you would find on Windows systems. If you are afraid of losing your favorite desktop diversion (such as Solitaire, FreeCell, and Minesweeper) when you leave Windows, have no fear. You can find many of them under GNOME games.
Table 7-1 lists the games available by selecting Games from the Red Hat menu. In this release, many KDE games (shown in Table 7-2) are also on this menu.

Table 7-1: GNOME Games

Game	Description
AisleRiot (solitaire)	Lets you select from among 28 different solitaire card games.
Chess	Gnuchess game in X. (Runs the `xboard` and `gnuchess` commands.)
Chromium Configuration	Set options such as skill level, screen size, and sound for Chromium.
Chromium	Deliver supplies to troops in battle in this action game.
FreeCell	A popular solitaire card game.
Freeciv (Isometric tileset)	In this strategy game, you try to lead your civilization to extinguish all others. (Uses Isometric tileset to represent cities, oceans, and other terrain.)
Freeciv Server (new game)	Server program needed to play Freeciv.
Gataxx	Board game where you flip over circles to consume enemy pieces.
Glines	Match five colored balls in a row to score points.
Gnect	Drop balls to beat the game at making four in a row.
Gnibbles	Steer a worm around the screen while avoiding walls.
GNOME Robots	Later version of Gnobots, which includes movable junk heaps.
GNOME Mines	Minesweeper clone. Click on safe spaces and avoid the bombs.
GNOME Stones	Move around a cave, collect diamonds, and avoid rocks.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 7-1: GNOME Games

Game	Description
GNOME Tetravex	Tetravex clone. Move blocks so that numbers on each side align.
Gnotski	Move pieces around to allow one piece to escape.
GTali	Yahtzee clone. Roll dice to fill in categories.
Iagno	Flip black and white chips to maneuver past the opponent.
Maelstrom	Navigate a spaceship through an asteroid field.
Mahjongg	Classic oriental tile game.
Same GNOME	Eliminate clusters of balls for high score.
Tux Racer	Steer a penguin as he races down a hill on his belly.

KDE games

If you install KDE, there are a bunch of games in the kdegames package. If you did not install the KDE desktop, you can install the kdegames package separately from CD #2. To see the KDE games (along with some GNOME games) on the Red Hat menu, select Games, then choose the game you want. The games available in KDE are listed by category in Table 7-2.

Table 7-2: Games for the KDE Desktop

Game	Description
Arcade Games	
KAsteroids	Destroy asteroids in the classic arcade game.
KBounce	Add walls to block in bouncing balls.
KFoul Eggs	Squish eggs in this Tetris-like game.
Klickety	Click color groups to erase blocks in this adaptation of Clickomania.
Kolf	Play a round of virtual golf.
KSirtet	Tetris clone. Try to fill in lines of blocks as they drop down.
KSmileTris	Tetris with smiley faces.
KSnakeRace	Race your snake around a maze.
KSpaceDuel	Fire at another spaceship as you spin around a planet.
KTron	Snake-style race game.
Boardgames	
Atlantik	Play this Monopoly-like game against other players on the network.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 7-2: Games for the KDE Desktop

Game	Description
KBackgammon	Online version of backgammon.
KBattleship	Sink the opponent's battleship in this online version of the board game.
KBlackBox	Find hidden balls by shooting rays.
Kenolaba	Move game pieces to push opponents' pieces off the board.
KMahjongg	Classic oriental tile game.
KReversi	Flip game pieces to outmaneuver the opponent.
Shisen-Sho	Tile game similar to Mahjongg.
Kwin4	Drop colored pieces to get 4-in-a-row.
Cardgames	
Patience	Choose from nine different solitaire card games.
KPoker	Video poker clone. Play five-card draw, choosing which cards to hold and which to throw
Lieutenant Skat	Play the card game Skat.
Megami	Play four blackjack hands against a dealer.
Tactics and Strategy	
KJumping Cube	Click squares to increase numbers and take over adjacent squares.
KAtomic	Move pieces to create different chemical compounds.
Konquest	Expand your interstellar empire in this multiplayer game.
Kolor Lines	Move marbles to form five-in-a-row and score points.
KMines	Minesweeper clone. Click safe spaces and avoid the bombs.
KSokoban	The Japanese warehouse keeper game.
SameGame	Erase game pieces to score points.

The games on the KDE menu range from diverting to quite challenging. If you are used to playing games in Windows, KMines and Patience will seem like old favorites. KAsteroids and KPoker are good for the mindless game category. For a mental challenge (it's harder than it looks), try KSokoban. For a challenging multiuser game on the GNOME menu, try Freeciv. For chess enthusiasts, there is Chess (xboard version of gnuchess).

The following sections describe a couple of the more interesting games that come with Red Hat Linux. First is the xboard game and some related chess programs. Next is a description of Freeciv.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chess games

Chess was one of the first games played on computer systems. While the game hasn't changed over the years, the way it's played on computers has. The set of chess programs that come with Red Hat Linux lets you play against the computer (in text or graphical modes), have the computer play against itself, or replay stored chess games. You can even play chess against other users on the Internet using Internet Chess Servers (ICS).

The `xboard` program is an X-based chess game that provides a graphical interface for gnuchess. GNU Chess (represented by the gnuchess package) describes itself as a communal chess program. It has had many contributors, and it seeks to advance a "more open and friendly environment of sharing" among the chess community. With `xboard`, you can move graphical pieces with your mouse. To play against the computer, click Games → More Games →Chess from the Red Hat menu, then start by just moving a piece with your mouse. While in the `xboard` window, select Mode → Two Machines to have the computer play itself. Select File → Load Game to load a game in Portable Game Notation (PGN). Figure 7-1 shows the `xboard` window with a "Two Machines" game in progress.

Figure 7-1: In the xboard window, you can set xgame to either play against the computer or to replay saved games.

You can use xboard to play online against others by connecting an xboard session to an Internet Chess Server (ICS). To start `xboard` as an interface to an ICS, either click Extras → Games → Chess or type the following command line:

$ **xboard -ics -icshost *name***

In this example, *name* should be replaced by the name of the ICS host (see the list of hosts below). In ICS mode, you can just watch games, play against other users, or replay games that have finished. The ICS host acts as a gathering place for enthusiasts who want to play chess against others on the Internet, watch games, participate in tournaments, or just meet chess people. Here are ICS host computers you can connect to by substituting *name* above with the address shown below:

Internet Chess Club: ICC (`chessclub.com`)
Chess.net (`chess.net`)
Free Internet Chess Server (`freechess.org`)

Freeciv

With Freeciv, you create a civilization that challenges competing civilizations for world dominance. The version of Freeciv that comes with Red Hat Linux contains both client software (to play the game) and server software (to connect players together). You can connect to your server and try the game yourself or (with a network connection) play against up to 14 others on the Internet.

You can start Freeciv from the Red Hat menu (as a non-root user) by clicking Games → More Games → FreeCiv (Isometric tile set). If Freeciv doesn't start, try starting it from a Terminal window by typing:

$ **civ &**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Figure 7-2 shows the two windows that appear when you start Freeciv. The Connect to Freeciv Server window contains your user name, host name, and port number. The Freeciv window is where you play the game.

Figure 7-2: Play Freeciv to build civilizations and compete against others.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If Freeciv won't start, one reason may be because you are logged in as root. You must be logged in as a regular user to run the civ command.

Starting Freeciv

You can play a few games by yourself, if you like, to get to know the game before you play against others on the network. The following procedure describes how to start your first practice Freeciv game.

1. Start Freeciv from the Red Hat menu (Games → More Games → FreeCiv, or

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

type **civ&**).
The Freeciv windows appear, as shown in Figure 7-2.

2. From a Terminal window, start the Freeciv server by typing:
3.
4. **$ civserver**
5. This is the server for Freeciv version 1.13.0
6. You can learn a lot about Freeciv at http://www.freeciv.org/
7. 2: Now accepting new client connections.
8.
9. For introductory help, type 'help'.
   >
10. Click Connect from the Connect to Freeciv Server window.
11. From the server prompt, type the following:
12. > **start**
13. Starting game.
14. 2: Loading rulesets
    >
    A "What Nation Will You Be?" window appears on the client, as shown in Figure 7-3.

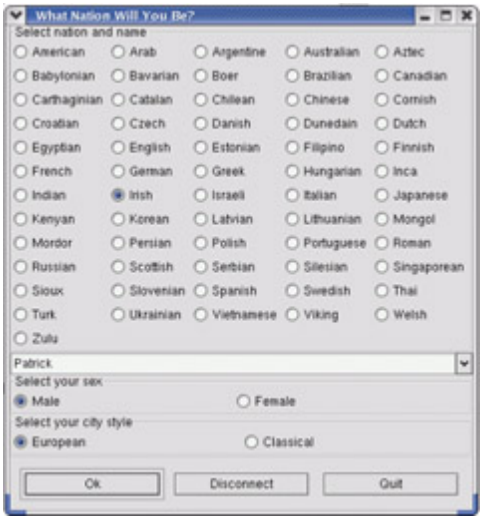

Figure 7-3: Choose a nation to begin Freeciv.

15. After you start Freeciv from the server prompt, choose a nation, the name of a leader, your sex, and the style of the city, and then click OK. At this point, you are ready to return to the Freeciv window.

Beginning with Freeciv

Check out the Freeciv window. Here are things you should know when you are starting. (You can find more help at the Freeciv site: `www.freeciv.org`.)

Click the Help button for topical information on many different subjects that will be useful to you as you play.

The world (by default) is 80 x 50 squares, with 11 x 8 squares visible at a time. The active square contains an icon of the active unit (flashing alternatively with the square's terrain).

Some squares contain special resources. Press and hold the middle mouse button for information on what special resources a square contains. (With a two-button mouse, hold the Ctrl key and click the right mouse button.) Try this a few times to get a feel for the land around you. This action also identifies any units on the terrain, as well as statistics for the unit.

To see the world outside of your 11 x 8 viewing area, click the scroll bars outside of the map. At first the world outside will be black. As units are added, areas closer to those units will be visible. (Press the letter **c** to return to the active part of your map.)

An overview map is in the upper-left corner of the Freeciv window. As the world becomes more civilized, this provides a good way to get an overview of what is going on. Right-click a spot on the overview map to have your viewport centered there.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The menu bar contains buttons you can use to play the game. The Game menu lets you change settings and options, view player data, view messages, and clear your log. The Kingdom menu lets you change tax rates, find cities, and start revolutions. The View menu lets you place a grid on the map or center the view. The Orders menu is where you choose the items you build and the actions you take. The Reports menu lets you display reports related to cities, military, trade, and science, as well as other special reports.
A summary of the economy of your civilization appears under the overview map. Information includes number of people, current year, money in the treasury, and percent of money distributed to tax, luxury, and science.
Ten icons below the overview information represent how money is divided between luxuries (an entertainer), research (a researcher), and taxes (a tax collector). Essentially, these icons represent how much of your resources are placed into improving each of those attributes of your community.
When you have made all your moves for a turn, click Turn Done. Next to that, a lightbulb indicates the progress of your research (increasing at each turn). A sun icon starts clear, but becomes brighter from pollution to warn of possible global warming. A government symbol indicates that you begin with a despotic government. The last icon tells you how much time is left in a turn.

The Unit box shows information about your current unit. You begin with two Settlers units and one Explorer unit.

Building your civilization

Start building your civilization. Here are things to try, as suggested by the Freeciv manual:

To change the distribution of money, choose Kingdom → Tax Rates. Move the slider bars to redistribute the percentage of assets assigned to luxury, research, and taxes. Try increasing research and reducing taxes to start off.
Change the current unit to be a settler as follows: click the stack of units on the map and click one of the Settlers from the menu that appears.
Begin building a city by clicking on Orders → Build City. When prompted, type a name for the city and click OK. The window that appears shows information about the city. It starts with one happy citizen, represented by a single icon (more citizens will appear as the game progresses).
The Food, Prod, and Trade lines reflect the raw productivity statistics for the city. The first number shows how much is being produced, the second (in parens) shows the surplus above what is needed to support the units. The Gold, Luxury, and Science lines indicate the city's trade output. Granary numbers show how much food is stored and the size of the food store. The pollution level begins at zero.
The Units at this point are not yet supported by a city (so nothing appears under Supported Units). When Units require support, they will be assigned to cities, and they will draw on city resources. Units present appear under that heading.
The map area shown consists of 21 squares that make up the city. The number 1 indicates the size of the city. The number 211 reflects the production of food, manufacturing production, and trade, respectively. The number 210 shows where the city's citizen is working and the results of the work.
The Phalanx line shows that the city can build a Phalanx and that it will take 20 production points to produce. Click Change to view other units the city could produce, select one you want to build, and click Change. Below that is a list of your current buildings (of which you have only a Palace to start out).
Close the city window by clicking on Close.

Exploring your world

To begin exploring, move the Settler.

1. Using the numeric keypad, press the 9 key three times to begin exploring. You can move the explorer up to three times per turn. You begin to see more of the world.
2. When the next unit (the Settler) begins blinking, move it one square in another

direction. Click Turn Done. Information for the city will be updated.

3. Click the City to see the city window. Notice that information about the city has been updated. In particular, you should see food storage increase. Close the city window.
4. Continue exploring and build a road. With the explorer flashing, use the numeric keypad to move it another three sections. When the settler begins blinking, press **r** to build a road. A small R appears on the square to remind you that the Settler is busy building a road. Click Turn Done.

Using more controls and actions

Now that you have some understanding of the controls and actions, the game can begin taking a lot of different directions. Here are a few things that may happen next and things you can do:

After you take a turn, the computer gets a chance to play as well. As it plays, its actions are reported to you. You can make decisions on what to do about those actions. Choose Game → Message Options. The Message options window appears, containing a listing of different kinds of messages that can come from the server and how they will be presented to you.

As you explore, you will run into other explorers and eventually other civilizations. Continue exploring by selecting different directions on your numeric keypad.

Continue to move the Settler one square at a time, once it has finished creating the road. (The Settler will blink again when it is available.) Click Turn Done.

At this point, you should see a message that your city has finished building Warriors. When buildings and units are complete, you should usually check out what has happened. Click the message associated with the city, then click Popup City. The city window appears, showing you that it has additional population. The food storage may appear empty, but the new citizens are working to increase the food and trade. You may see an additional warrior unit. A science advisory may also appear at this point to let you choose your city's research goals. Click Change and select Writing as your new research goal. You can then select a different long-term goal as well. Click Close when you are done.

If your new Warrior is now blinking, press **s** to assign sentry mode to the Warrior.

You should be familiar with some of the actions of Freeciv at this point. To learn some basic strategies for playing the game, choose Help → Help Playing.

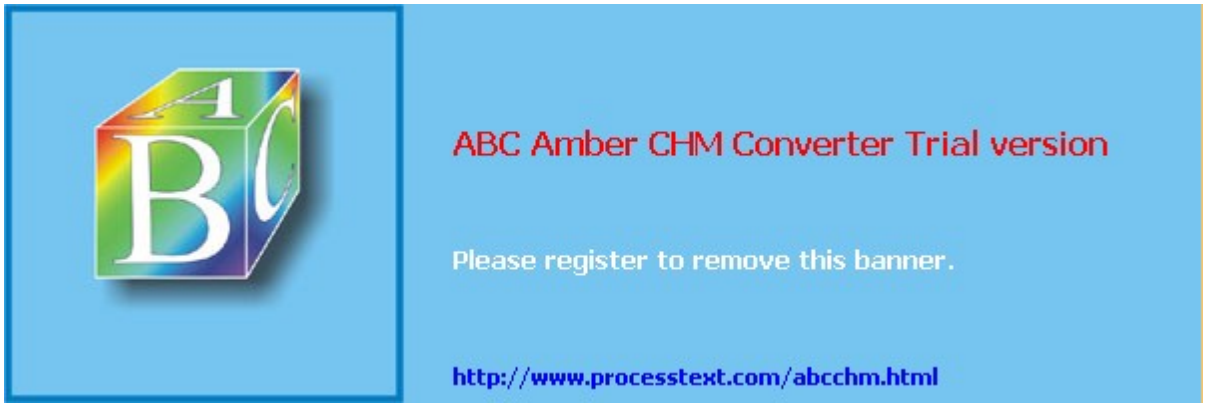

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Commercial Linux Games

When Loki Software, Inc. closed its doors a few years ago, the landscape of commercial gaming in Linux changed. Loki produced Linux ports of popular games, including Myth II and Civilization: Call to Power, to name a few. Today, commercial games that run natively are lead by several popular games from id Software (described in the next section).

In the wake of Loki's demise, TransGaming Technologies has been working on an approach to bringing popular games to Linux that relies on a version of WINE called WineX. In most cases, instead of having different ports of popular games (as Loki did), Transgaming lets users run existing Windows games in Linux by adapting WineX to each game that needs a tweak here and there.

Although Loki Software, Inc. is gone, certain Loki Games are still available for purchase on the Web. Although they sell for a fraction of their original price, you are on your own if they don't work since Loki Software is no longer there to support them. The Loki Games Demo is still around, if you want to get a feel for a particular Loki games before it disappears completely (I describe how to find demo and packaged Loki Games later in this chapter).

id Software Games

Among the most popular games running natively in Linux are Quake III Arena, Return to Castle Wolfenstein, and Unreal Tournament from id Software, Inc. You can purchase Linux versions of these games or download demos of each game before you buy.

Note If you have trouble getting any id Software games running in Linux, refer to the Linux FAQs available from id Software at: http://zerowing.idsoftware.com/linux.

Quake III Arena

Quake III Arena is a first-person shooter-type game where you can choose from lots of weapons (lightning guns, shotguns, grenade launchers, and so on) and pass through scenes with highly detailed 3D surfaces. You can play alone or against your friends. There are multiplayer death-match and capture-the-flag competitions.
A demo version of Quake III Arena for Linux is available from id Software's FTP site (`ftp://ftp.idsoftware.com/idstuff/quake3/linux`). Look for the file `linuxq3ademo-1.11-6.x86.gz.sh`. Figure 7-4 shows a screenshot from Quake III Arena.

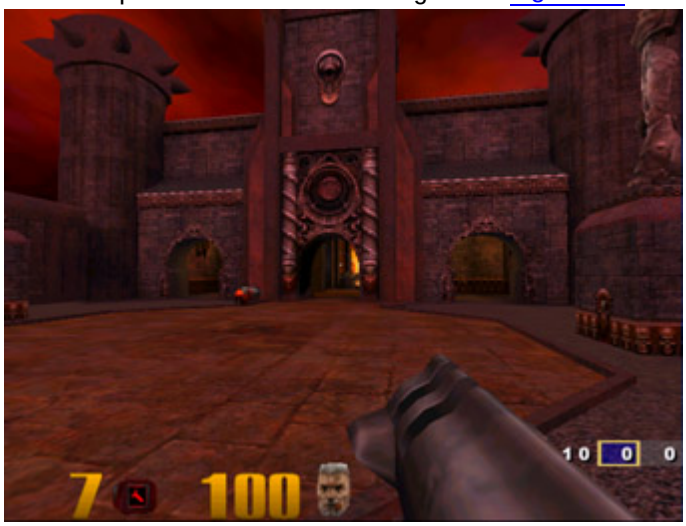

Figure 7-4: Quake III Arena is a popular first-person shooter game that runs in Linux.

Unreal Tournament

In Unreal Tournament, futuristic gladiators with super-human strength and advanced weaponry battle each other to find the ultimate warriors. The latest version is Unreal Tournament 2003. You can choose single-player or multi-player games where you compete by yourself or in teams.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Warning Unfortunately, Unreal Tournament 2003 requires an OpenGL extension that is currently only available in NVIDIA cards (GL_EXT_texture_compression_s3tc). If you have an NVIDIA card, be sure to download the binary-only Linux driver for your NVIDIA card before running UT 2003.

A demo version of Unreal Tournament 2003 for Linux is available from id Software. As with the full version, the demo requires you to have an NVIDIA video card. You can get the demo from the following downloads page: `www.unrealtournament2003.com/?downloads`. Figure 7-5 shows the first screen of the Unreal Tournament 2003 demo for Linux:

Figure 7-5: Battle against super-human gladiators with Unreal Tournament 2003.

Return to Castle Wolfenstein

Mixing World War II action with creatures conjured up by Nazi scientists, you battle with the Allies to destroy the Third Reich. Return to Castle Wolfenstein is based on the Quake III Arena engine. The game offers single-player mode as well as team-based multiplayer mode.

If you purchase Return to Castle Wolfenstein for Linux, you actually get the Windows version with an extra Linux installer. If you already have the Windows version, you can download the Linux installer and follow some instructions to get it going. I downloaded the installer called `wolf-linux-1.33.x86.run` from `ftp://ftp.idsoftware.com/idstuff/wolf/linux`. The INSTALL file (in `/usr/local/games/wolfenstein`) describes what files you need to copy from the Windows CD.

To get a demo of Return to Castle Wolfenstein, visit `ftp.idsoftware.com/idstuff/wolf/linux/old`. The demo I used was a single-player demo (`wolfspdemo-linux-1.1b.x86.run`). A multiplayer demo is also available.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Warning As with Unreal Tournament 2003, you need an NVIDIA card to run Return to Castle Wolfenstein.

Figure 7-6 is a screenshot from Return to Castle Wolfenstein running in Linux.

Figure 7-6: Return to Castle Wolfenstein combines strange creatures and WWII battles.

TransGaming and WineX gaming

TransGaming Technologies brings to Linux some of the most popular games that currently run on the Windows platforms. Working with WINE developers, TransGaming is developing WineX, which enables you to run many different games on Linux that were originally developed for Windows. Although TransGaming is producing a few games that are packaged separately and tuned for Linux, in most cases it sells you a subscription service to WineX instead of the games. That subscription service lets you stay up-to-date on the continuing development of WineX so you can run more and more Windows games.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Unfortunately, the kernel that ships with Red Hat Linux does not work with WineX. To use WineX in Red Hat Linux, TransGaming recommends that you download a vanilla kernel from

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To get Windows games to run in Linux, WineX particularly needs to develop Microsoft DirectX features that are required by many of today's games. There are also issues relating to CD keys and hooks into the Windows operating system that must be overcome (such as requiring Microsoft Active Desktop). In fact, a WineX subscription has value, in part, because it lets you vote on which games you'd like to see TransGaming work on next.

A full list of games supported by TransGaming, as well as indications of how popular they are and how well they work, is available from the TransGaming site (search for the games that interest you by going to www.transgaming.com and clicking on Games). More than 100 games are currently listed with a rating of 4 out of 5 (meaning that the game will run well, if not flawlessly). One game, The Sims, which is available in a separate, optimized Linux version from TransGaming, is rated a 5 (meaning that it runs flawlessly).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The Windows version of Sims doesn't run under WineX. You need to purchase Sims for Linux. Although TransGaming believes that Sims for Linux will run on any recent Linux distrib

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Support or major enhancements for the following games was recently added to WineX 3.1:

- Morrowind (new support)
- Grand Theft Auto: Vice City (new support)
- Battlefield 1942 (graphical and speed improvements)
- Medal of Honor: Allied Assault (corrected defect that caused game to crash)

Here is a list of some of the most popular games that are being used by the TransGaming community (though they may or may not have been tested by TransGaming) and that are said to run well in WineX (rated a 4 or 5). I suggest you check the TransGaming list yourself for additions and changes:

- Half-Life and Counter-Strike
- Diablo II
- StarCraft
- Baldur's Gate 2
- Jedi Knight 2: Jedi Outcast
- Max Payne
- Return to Castle Wolfenstein
- American McGee's Alice
- Command & Conquer Red Alert 2
- Baldur's Gate
- Total Annihilation
- DeusEx
- The Sims (Mandrake Gaming Edition)
- Fallout 2
- Warcraft II

With WineX 3.1, Transgaming has added a new Point2Play feature. Point2Play provides a graphical window for installing, configuring, and testing WineX on your computer. It also lets you install and organize your games so you can launch them graphically. Figure 7-7 shows an example of the Transgaming Point2Play window:

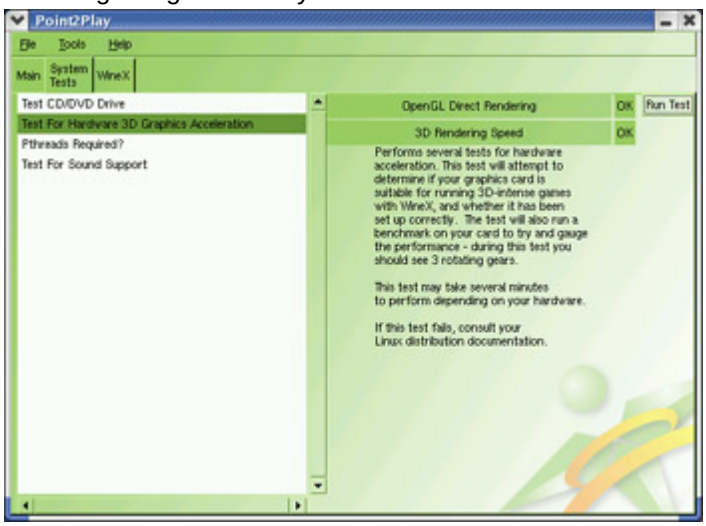

Figure 7-7: Use the Point2Play window to check computer hardware for WineX gaming.

Other features in the new Point2Play window include the ability to select among different installed versions of WineX for running applications and tools for individually configuring how each game runs under WineX. (If a game won't run from the GUI, try launching it from a Terminal window.

To get binary copies (ones that are already compiled to run) of WineX and Point2Play, you need to subscribe to TransGaming. For details on how to become a "TransGamer," click on the Subscribe Here link on the TransGaming home page (`www.transgaming.com`). Benefits currently include:

- Downloads of the latest version of WineX.
- Access to WineX support forums.
- Ability to vote on which games you want TransGaming to support next.
- Subscription to the WineX newsletter.

Source code for WineX is should be available in the near future if you want to build your own WineX package. To check availability, try the SourceForge.net project site for WineX (`sourceforge.net/projects/winex`).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Loki Software game demos

To encourage people to get to know their games, Loki offered a demo program that let you choose from among more than a dozen of its games to download and try. Although Loki Software, Inc. is no longer in business, you can still find some of its games for sale. For example, a recent search for "Loki" at Amazon.com turned up 16 different Loki games (including the ones described here), with many selling for $9.99.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Warning If you try to download any of the demos described in the next sections, make sure you have plenty of disk space available. It is common for one of these demos to require

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The Loki Demo Launcher for downloading demos is still available from the Demo Launcher page (www.lokigames.com/products/demos.php3). From that page, there are links to FTP sites from which you can download the Demo Launcher. (If that page is gone, try ftp://david.hedbor.org/loki/demos.) The file that you want to save is `loki_demos-full-1.0e-x86.run`. Save it to a directory (such as `/tmp/loki`) and do the following:

1. Change to the directory where you downloaded the demo. For example:
2. # **cd /tmp/loki**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note You may need to be root user to install these games. However, the default paths where the DemoLauncher tries to write by default are only accessible to the root user.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

3. As root user, run the following command (the program may have a different name if it has been updated):
4. # **sh loki_demos-full-1.0e.x86.run**

   If you have not used the Demo Launcher before, a screen appears asking you to identify the paths used to place the Install Tool.
5. If the default locations shown are okay with you, click Begin Install.

   Assuming that there was no problem writing to the install directories, you should see an Install Complete message.
6. Click Exit. Next, you should see the Uninstall Tool window.
7. If the paths for holding the Uninstall Tool are okay, click Begin Install. The Install Complete message appears.
8. Click Exit. The window that appears allows you to set the locations for installing the Demo Pack.
9. If the paths are okay, click Begin Install.

   Next, you should see a box that shows the different demo games that are available. As you move the cursor over each game, the disk space is displayed for that game. Click the games you want to install and then click Continue. A window appears, displaying the progress of each download.
10. You may need to click an Update button to complete the update and Finish to finish it.

    The demo should now be ready to start.
11. Either click Play or type `loki_demos` from a Terminal window to start the program.
12. Select to start the game, and you're ready to go.

The following sections describe a few games that may still be available. Again, these games may not be available for long.

Civilization: Call to Power

You can build online civilizations with Civilization: Call to Power (CCP). Like earlier versions and public spin-offs (such as the Freeciv described earlier in this chapter), Civilization: Call to Power for Linux lets you explore the world, build cities, and manage your empire. This latest version offers multiplayer network competition and extensions that let you extend cities into outer space and under the sea.

If you like the Freeciv game that comes with Red Hat Linux, you will love CCP. Engaging game play is improved with enhanced graphics, sound, and animation. English, French, German, Italian, and Spanish versions are available.
The CCP demo comes with an excellent tutorial to start you out. If you have never played a civilization game before, the tutorial is a great way to start. Figure 7-8 shows an example of a scene from the Civilization: Call to Power for Linux demo.

Figure 7-8: Civilization Call to Power features excellent graphics and network play.

Myth II: Soulblighter

If you like knights and dwarves and storming castles, Myth II: Soulblighter for Linux might be for you. In Myth II, you are given a mission and some troops with various skills. From there, you need strategy and the desire to shed lots of virtual blood to meet your goal.
Myth II was created by Bungie Software (`www.bungie.com`) and ported to Linux by Loki Entertainment Software (`www.lokigames.com`). This version of the popular Myth game includes improved graphics

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

and new scenarios.

A demo version is available that runs well in Red Hat Linux. You can get it via the Demo Launcher described earlier. As usual, you will need a fairly powerful computer (at least a Pentium 133 MHz, 32MB RAM, 80MB swap space, and 100MB of free disk space). You need network hardware for multiuser network play (network card or dialup) and a sound card if you want audio. A screenshot of Myth II is shown in Figure 7-9.

Figure 7-9: Use warriors, archers, and dwarves to battle in Myth II.

Heretic II

Based on the Quake Engine, Heretic II sets you on a path to rid the world of a deadly, magical plague. As the main character, Corvus, you explore dungeons, swamps, and cities to uncover and stop the plague. The graphics are rich and the game play is quite engaging.

You will experience some crashing problems with Heretic II out-of-the-box. Be sure to check for the update to Heretic II at `updates.lokigames.com`, which should fix most of the problems.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version

Please register to remove this banner.

http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

While Linux has not yet become a dominant gaming platform, you can still spend your time on plenty of games running on Red Hat Linux. Old UNIX games that have made their way to Linux include a variety of X Window-based games. There are card games, strategy games, and some action games.

On the commercial front, Civilization: Call to Power for Linux and Myth II are available to use on your Red Hat Linux system, but will soon disappear because Loki Software (which ported those applications to Linux) went out of business. The future of high-end Linux gaming might be in the hands of TransGaming Technologies, which has created WineX from current WINE technology to allow Windows games to run in Linux.

Commercial games that run natively in Linux are also available. These include games from id Software, such as Quake III Arena, Unreal Tournament, and Return to Castle Wolfenstein.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 8: Multimedia in Red Hat Linux

Overview

In this chapter, you learn to add sound, video, digital images, and other multimedia support to your Red Hat Linux system. The chapter describes how to configure a sound card to work with Red Hat Linux. Then it tells you how to get and use software to play music CDs and a variety of sound formats from your Red Hat Linux system.

Video clips that are readily available on the Internet in formats such as AVI, OGG, and QuickTime can be viewed using several different players in Red Hat Linux. Also, you can view live television and video using TV cards and Webcams.

Because CD-ROM is the medium of choice for recorded music, this chapter describes how to set up and use CD burners to create your own music CDs. Once your CD burner is set up to record music, the same CD burner can be used to back up your data or to create software CDs.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Listening to Audio

If you want the sound on your computer to be more than the speaker on your PC going "bing" occasionally, you must have a separate sound card or one built into your computer's motherboard. Games are a reason to add a sound card, if the games rely on sound effects or other audio cues. Also, sound cards let you play music and communicate on the Internet using a variety of communications tools.

To give you an idea of the features that a sound card can provide, the following lists features that are included in the popular Sound Blaster 16 sound card:

Sound recording and playback — The card can convert analog sound into 8-bit or 16-bit digital numbers. To convert the sound, the board samples the sound in waves from 5 kHz to 48 kHz, or 5000 to 48,100 times per second. (Of course, the higher the sampling, the better the sound and larger the output.)

Full-duplex support — This allows for recording and playback to occur at the same time. This is particularly useful for bidirectional Internet communication or simultaneous recording and playback.

Input/output ports — Several different ports on the board enable you to connect other input/output devices. These ports include:

Line-In — Connects an external CD player, cassette deck, synthesizer, MiniDisc, or other device for recording or playback. If you have a television card, you might also patch that card's line out to your sound card's line in.

Microphone — Connects a microphone for audio recording or communications.

Line-Out (Speaker Out) — Connects non-powered speakers, headphones, or a stereo amplifier.

Joystick/MIDI — Connects a joystick for gaming or MIDI device.

Internal CD Audio — This internal port connects the sound card to your computer's internal CD-ROM board (this port isn't exposed when the board is installed).

Sound drivers provided in Linux come from many sources, including a project that no longer exists: Open Sound System/Free (OSS/Free). If your sound card isn't supported in Red Hat Linux, try the OSS/Linux package to see if that package supports your sound card (go to `www.opensound.com/oss.html`). If it works, you can purchase the sound package.

Caution Before you install a separate sound driver distribution, check to see if your current Red Hat distribution already has the most recent sound driver. When possible, use the driver that's distributed with the kernel. If you have tried the procedures in this book and you still don't have a working sound card, read the following file: /usr/src/linux*/Documentation/sound/README.OSS. Also, you can visit www.linux.org.uk/OSS for a list of supported cards. Or, you could try the Advanced Linux Sound Architecture at www.alsa-project.org, which is another project offering open-source sound drivers.

The devices that the audio programs use to access audio hardware in Red Hat Linux include:

/dev/audio, /dev/audio1 — Devices that are compatible with Sun workstation audio implementations (audio files with the .au extension).

/dev/cdrom — Device representing your first CD-ROM drive. (Additional CD-ROM drives are located at `/dev/cdrom1`, `/dev/cdrom2`, etc.)

/dev/dsp, /dev/dsp1 — Digital sampling devices, to which many audio applications direct sound.

/dev/mixer, /dev/mixer1 — Sound-mixing devices.

/dev/sequencer — Device that provides low-level interface to MIDI, FM, and GUS.

/dev/midi00 — Device that provides raw access to midi ports.

/dev/sndstat — Device that displays the status of sound drivers.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Nodes in the /dev directory, such as /dev/audio, aren't just regular files. They represent access points to the physical devices (hard disks, COM ports, etc.) that are connected to your system, or to pseudo-devices (such as Terminal windows). For example, to find out the device of your current Terminal window, type **tty**. Then send some data to that device. For example, if your device name is /dev/pts/0, type:
$ **echo "Hello There" > /dev/pts/0**

The words "Hello There" appear in that Terminal window. You can try sending messages among several Terminal windows. If a user who is logged on to the computer has terminal permissions open, you can send messages to him or her in this way, too. (I knew people who would send a dictionary file to an unsuspecting user's terminal. Although it wasn't destructive, it was annoying if you were trying to get work done.)

For general information about sound in Linux, see the Sound-HOWTO (for tips about sound cards and general sound issues) and the Sound-Playing-HOWTO (for tips on software for playing different types of audio files).

Note You can find Linux HOWTOs at www.tldp.org.

Configuring your sound card

During the first start-up after you install Red Hat Linux, the Red Hat Setup Agent tries to detect and configure your sound card. If that process was successful, you can skip this procedure. If your sound card wasn't detected or if you add a card later, here are a few things you can try:

Insert a music CD into the drive and see if it plays.

Tip Audio volume is set fairly low by default. If you can barely hear your audio, click Sound & Video → Volume Control (from the Red Hat menu) to increase the volume level. Move the sliders to adjust the volume to your taste. Alternatively, from the command line type **aumix -v 100** in a Terminal window to turn the volume all the way up, so you can adjust sound from your volume knobs.

From the Red Hat menu, click System Settings → Soundcard Detection (this runs the `redhat-config-soundcard` command). If your sound card was detected, an Audio Devices window should appear, as shown in Figure 8-1. Click the Play Test Sound button and you should hear a test sound. Click OK to finish up.

Figure 8-1: The Audio Devices (redhat-config-soundcard) window detects your sound card.

If the sound card still isn't detected, try the Sound Configuration Utility (you may need to install the sndconfig package from CD #3), as follows:

1. Run the following command as root user from a Terminal window:
2. # **/usr/sbin/sndconfig**
3. A sndconfig window is displayed, asking ask if you want to probe for Plug-and-Play cards. Press Enter to select OK.
4. If the utility finds a sound card that it recognizes, it displays the name of that card and asks if you want to configure it. If the sound card was detected, press Enter to select OK.
5. A message warns you that you are replacing the old `/etc/conf.modules` file.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Press Enter to select OK.

6. A message tells you that a sound sample will play to check if the card is configured properly. Press Enter to select OK.
7. You are asked if you heard the sound.
8. If you hear the sound, select Yes and skip to Step 12. Otherwise select No to continue.
9. You are asked to manually configure your sound card. Press Enter to select OK.
10. A list of sound cards is displayed. Select your card and press Enter to select OK.
11. Select OK to try the sound card again.
12. You are asked again if you were able to hear the sample. If you hear the sound, select Yes to continue. Otherwise select No to try a different card.
13. You are warned that a new `/etc/modules.conf` file is being created (and the old one is being backed up). Press Enter to select OK.

At this point, you can try playing some audio. Insert a CD and open one of the CD players described in the following section.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If there is a data CD in your CD drive, you may not be able to simply eject it to play your music CD. To eject a data CD, close any windows that may have an open file from

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Choosing an audio CD player

The player that pops up automatically when you insert a CD (on the GNOME desktop) is the gnome-cd player. However, several CD players that come with Red Hat Linux can be used to play CDs. Here are your choices for playing CDs with Red Hat Linux:

CD Player (`gnome-cd`) — This is the default CD player for the GNOME desktop. Besides having standard play buttons, this player lets you get track information automatically from a CD database (such as `freedb.org`). Or, if your CD isn't listed in the database, you can enter your own track information manually.

KsCD Player (`kscd`) — The KsCD player comes with the KDE desktop. To use KsCD, the kdemultimedia package must be installed. From the Red Hat menu (KDE desktop), click Sound & Video → KsCD (or type **`kscd`** from a Terminal window). Like gnome-cd, this player lets you get title, track, and artist information from the CD database. KsCD, however, also lets you submit information you type in yourself to a CD database (if your CD isn't found there).

Grip (`grip`) — While the Grip window is primarily used as a CD ripper, it can also play CDs. Click Sound & Video → More Sound & Video Applications → Grip. It includes tools for gathering data from and submitting data to CD databases. It also includes tools for copying (ripping) CD tracks and converting them to different formats (encoding). (The grip package must be installed to use this command.)

CDPlay (`cdp`) — If you don't have access to the desktop, you can use the text-based `cdp` command. This player lets you use keyboard keys to play your CD, select tracks, go forward or back, or eject. (The cdp package must be installed to use this command.)

X Multimedia System (`xmms`) — The xmms player plays a variety of audio formats, but can also play directly from a CD.

If you prefer to use one of these players, you can disable automatic startup of gnome-cd by disabling automatic play in your CD Properties window (as described in the next section). The gnome-cd and cdp CD players are then described.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you try some of these CD players and your CD-ROM drive is not working, see the sidebar "Troubleshooting Your CD-ROM" for further information.

Automatically playing CDs

When you put an audio CD into your CD-ROM drive, a CD player automatically pops up on your desktop. If you are using the GNOME desktop, you can thank magicdev. The magicdev process monitors your CD-ROM drives and opens a CD player when it sees an audio CD.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The fact that inserting a CD starts a player automatically is nice to some people and annoying to others. If you just want the CD to play, this behavior is a good thing. However, if you want to choose your own CD player or not play the CD until you choose, you may find auto-playing a bother. If you insert a data CD or a blank CD, magicdev exhibits different behavior. Here is what magicdev does by default:

Audio CD — When the music CD is inserted, magicdev starts the gnome-cd CD player and starts to play the first track of the CD.
Data CD — When a data CD is inserted, the CD is mounted on your file system, any auto-run program that may be on the CD is launched, and a file manager window opens. The first CD drive's mount point (`/dev/cdrom`) is `/mnt/cdrom`. If you have two drives, the second (`/dev/cdrom1`) is mounted on `/mnt/cdrom1` (and so on).
Blank CD — When a blank CD is inserted, a nautilus window opens with `burn:///` as the location.

You can change the behavior of magicdev on the CD & DVD preferences window. As you might guess from its name, you can change behavior for both CDs and DVDs as follows:

1. From the Red Hat menu, click Preferences → CD Properties. The CD & DVD preferences window appears.
2. For Data CDs, select from the following options:
   Mount CD when inserted — If this is selected, when a data CD is inserted it is automatically mounted in a subdirectory of `/mnt`. This option is on by default.
   Start auto-run program on newly mounted CD — If this is selected, after a data CD is mounted, the user is asked to choose whether to run an auto-run program from the CD. This option is on by default if the first option is checked.
   Open file manager window for newly mounted CD — If this is selected, after a data CD is mounted, a file manager window opens to display the contents of the top-level directory of the CD. For the first CD drive, the directory is `/mnt/cdrom`. This option is on by default if the first option is checked.
3. For Audio CDs, you can select "Run command when CD is inserted" to have the CD start playing automatically after it's inserted. The command shown in the box labeled Command is used to play the CD. By default, the option is on, and the gnome-cd player is chosen for you.
4. For blank CDs, the nautilus window opens with `burn:///` as the location. With this feature enabled, you can drag-and-drop files on the nautilus window to gather the files you want to write to CD. Click Write to CD to burn the selected files to the CD.
5. For DVDs, click the box next to "Run command when DVD (video) is inserted" to have the DVD play automatically (using the `vlc` command) when you insert a DVD.
6. Click Close.

Troubleshooting Your CD-ROM

If you are unable to play CDs on your CD-ROM drive, here are a few things you can check to correct the problem:

Check that your sound card is installed and working properly (see "Configuring your sound card" earlier in this chapter).
Check that the wire between your CD-ROM and sound card is connected.
Check that the CD-ROM was detected when you booted Linux. If your CD-ROM drive is an IDE drive, type **dmesg | grep ^hd**. You should see messages about your CD-ROM that look like this: `"hdc: CD-ROM CDU701, ATAPI CDROM drive"` or this: `"hdc: ATAPI 14X CD-ROM drive, 128kB Cache."`

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

If you see no indication of a CD-ROM drive, check that the power supply and cables to the CD-ROM are connected. To make sure that the hardware is working, you can also boot to DOS and try to access the CD.
Try inserting a software CD-ROM. If you are running the GNOME or KDE desktop, a window should appear indicating that the CD mounted by itself. If no such window appears, go to a Terminal window and type **mount /mnt/cdrom**. Then change to the `/mnt/cdrom` directory and list the contents with the command `cd /mnt/cdrom; ls`. This tells you if the CD-ROM is accessible.
If you get the CD-ROM working, but it fails with the message "CDROM device: Permission denied" when you try to play music as a nonroot user, the problem may be that `/dev/cdrom` (which is typically a link to the actual hardware device) is not readable by anyone but root. Type **ls -l /dev/cdrom** to see what the device is linked to. Then, as the root user, type **chmod 644 /dev/cdrom** to enable all users to read your CD-ROM and to enable the root user to write to it. One warning: if others use your computer, they will be able to read any CD you place in this drive.

Playing CDs with gnome-cd

Like most graphical CD players, the gnome-cd player has controls that look similar to what you would see on a physical CD player. If you are using the GNOME desktop, from the System Menu select Sound & Video → CD Player, or from a Terminal window, type:

$ **gnome-cd &**
The interface for adding information about the CD and its tracks is very nice. Click the Open Track Editor button. If you have an active connection to the Internet, `gnome-cd` automatically tries to grab track information for your CD, so you may already see the CD title and track name for each song when you open the editor.
If information about the CD doesn't appear, you can add Artist and Title information about the CD yourself. Then you can select each track to type in the track name. To add the name of the artist and the disc title, click in the appropriate box and type in that information. Figure 8-2 shows the CD Player and the CDDB Track Editor.

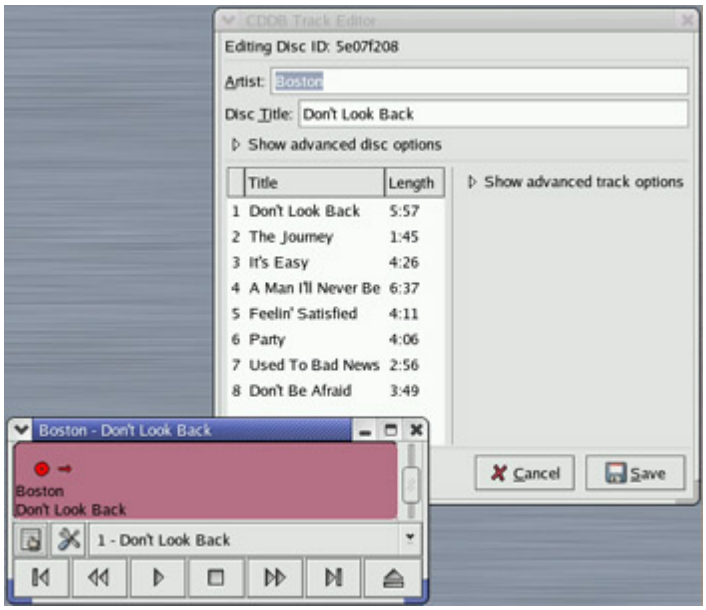

Figure 8-2: Play CDs and store artist, title, and track information with gnome-cd.

Playing CDs with cdp

If you are working from a dumb terminal or just don't have your X desktop running, you can run the `cdp` utility (which comes with Red Hat Linux) to play CDs. I don't suggest running this from a Terminal

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

window; it doesn't display properly. First, insert the music CD you want to play. Then, to start `cdp`, at a shell prompt type:

$ **cdp**

You should see a blue screen containing the `cdp` display. If instead of starting on the first track you want to start on another track (for example, track 5), type:

$ **cdp play 5**

When `cdp` starts, you can see all the tracks, how long each track plays, and total play time. To control the play of the CD, use the following controls (turn on Num Lock to use these numbers from the numeric keypad):

9 — Play
8 — Pause/Resume
7 — Stop
6 — Next Track
5 — Replay Current Track
4 — Previous Track
3 — Forward 15 Seconds
2 — Quit (Stop Music, Exit, and Eject)
1 — Back 15 Seconds
0 — Exit (Continue Music and Exit)
. — Help (Press the period key)

The `cdp` display also lets you enter the names of the artist, CD, and each song. Because this information is saved, you can see it each time you play the CD. Type these commands while the `cdp` display is showing to edit information about the CD currently playing:

a — Edit the Artist Name and press Enter.
c — Edit the CD Name and press Enter.
Enter — Edit the title of the current song and press Enter again.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution If you try to edit a song name and cdp crashes, type **eject** to stop the CD from playing. Editing the song name seems to work better if you pause the song first.

The arrow keys are also pretty handy for controlling CDs in `cdp`. The up arrow is for pause/play, and the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

left arrow is to go back a track. The right arrow is to go forward a track, and the down arrow is to eject.

Playing music with XMMS Audio Player

The XMMS (X Multimedia System) Audio Player provides a graphical interface for playing music files in MP3, Ogg Vorbis and WAV, and other audio formats. XMMS has some nice extras too, which include an equalizer, a playlist editor, and the ability to add more audio plugins. If the player looks familiar to you, that's because it is styled after the Windows winamp program.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Red Hat removed all software that does MP3 encoding or decoding due to patent concerns related to MP3 format. Although the XMMS player was designed to play MP3 files,

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You can start the XMMS Audio Player by selecting Sound & Video → Audio Player or by typing the `xmms` command from a Terminal window. Figure 8-3 consists of the XMMS Audio Player with the associated equalizer (below) and the Playlist Editor (to the right).

Figure 8-3: Play Ogg Vorbis and other audio files from the XMMS playlist.

As noted earlier, you can play several audio file formats. Supported audio file formats include the following:

- MP3 (with added plugin)
- Ogg Vorbis
- WAV
- AU
- CD Audio
- CIN Movies

You can get many more audio plugins from `xmms.org`. The XMMS Audio Player can be used in the following way:

1. Obtain music files by either:
   Ripping songs from a CD or copying them from the Web so that they are in an accessible directory.
   Inserting a music CD in your CD-ROM drive. (Xmms expects the CD to be accessible from `/dev/cdrom`.)
2. From the Red Hat menu, click Sound & Video → Audio Player. The X Multimedia System player appears.
3. Click the Eject button. The Load files window appears.
4. If you have inserted a CD, the contents of `/mnt/cdrom` appear in the Files pane. Select the files you want to add to your Playlist and click the Add Selected Files or click the "Add all Files in Directory" button to add all songs from the current directory. To add audio files from your file system, browse your files and directories and click the same buttons to add the audio files you want. Select Close.
5. Click the Play List button (the tiny button marked PL) on the console. A Playlist Editor window appears.
6. Double-click the music file, and it starts to play.
7. With a file selected and playing, here are a few actions you can take:
   Control play — Buttons for controlling play are what you may expect to see on a physical CD player. From left to right, the buttons let you go to a previous track, play, pause, stop, go to the next track, or eject. The eject button opens a window, allowing you to load the next file.
   Adjust sound — Use the left slider bar to adjust the volume. Use the right slider bar to change the right-to-left balance.
   Display time — Click in the elapsed time area to toggle between elapsed time and time remaining.
   View file information — Click the button in the upper-left corner of the screen to see the XMMS menu. Then select View File Info. You can often find out a lot of information about the file: title, artist, album, comments, and genre. For an Ogg file, you can see specific information about the file itself, such as the format, bit rate, sample rate, frames, file size, and more. You can change or add to the tag information and click Save to keep it.
8. When you are done playing music, click the Stop button to stop the current song. Then click the X in the upper-right corner of the display to close the window.

Special features of the XMMS Audio Player let you adjust high and low frequencies using a graphic equalizer and gather and play songs using a Playlist Editor. Click the button marked EQ next to the balance bar on the player to open the Equalizer. Click the button marked PL next to that to open the Playlist Editor.

Using the Equalizer

The Equalizer lets you use slider bars to set different levels to different frequencies played. Bars on the left adjust lower frequencies, and those on the right adjust higher frequencies. Click the EQ button to open the Equalizer. Here are tasks you can perform with the Equalizer:

If you like the settings you have for a particular song, you can save them as a Preset. Set each frequency as you like it and click the Preset button. Then choose Save → Preset. Type a name for the preset and click OK.

To reload a preset you created earlier, click the Preset button and select Load → Preset. Select the preset you want and click OK. The settings change to those of the preset you just loaded.

The small window in the center/top of the Equalizer shows the sound wave formed by your settings. You can adjust the Preamp bar on the left to boost different levels in the set range.

Using the Playlist Editor

The Playlist Editor lets you put together a list of audio files that you want to play. You can add and delete files from this list, save them to a file, and use them again later. Click the PL button in the XMMS window to open the Playlist Editor.

The Playlist Editor allows you to:

Add files to the playlist — Click the Add button. The Load Files window appears. Select the directory containing your audio files (it's useful to keep them all in one place) from the left column. Then either select a file from the right column and click Add selected files or click Add all files in the directory. Click OK. The selected file or files appear in the playlist.

Select files to play — To select from the files in the playlist, use the previous track and next track buttons in the main XMMS window. The selected file is highlighted. Click the Play button to play that file. Alternatively, you can double-click on any file in the playlist to start it playing.

Delete files from the playlist — To remove files from the playlist, select the file(s) you want to remove (next/previous track buttons), right-click the playlist window, and click Remove → Selected . The selected files are removed.

Sort files on the playlist — To sort the playlist in different ways, click and hold the Misc button and move the mouse to select Sort List. Then you can select Sort List to sort by Title, Filename, Path and Filename, or Date. You can also randomize or reverse the list.

Save the playlist — To save the current playlist, hold the mouse button down on the List button and then select Save. Browse to the directory you want, and then type the name you want to assign to the playlist and click OK.

Load the playlist — To reload a saved playlist, click the List button. Select a previously saved playlist from the directory in which you saved it and click OK.

There is also a tiny set of buttons on the bottom of the Playlist Editor screen. These are the same buttons as those on the main screen used for selecting different tracks or playing, pausing, stopping, or ejecting the current track.

Using MIDI audio players

MIDI stands for Musical Instrument Digital Interface. MIDI files are created from synthesizers and other electronic music devices. MIDI files tend to be smaller than other kinds of audio files because, instead of storing the complete sounds, they contain the notes played. The MIDI player reproduces the notes to sound like a huge variety of MIDI instruments.

There are lots of sites on the Internet for downloading MIDI files. Try the Ifni MIDI Music site (www.ifni.com), which contains songs by the Beatles, Led Zeppelin, Nirvana, and others organized by

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

album. Most of the MIDI music is pretty simple, but you can have some fun playing with it.
Red Hat Linux comes with a couple of MIDI players, including `kmid` and `kmidi`. Neither player is installed by default (they are both in the kmultimedia package on CD #3). Both provide a GUI interface for midi music. The more interesting of the two midi players is KMidi. There is also the timidity MIDI player, which lets you run MIDI audio from a Terminal window.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Use the timidity or kmidi MIDI player if your sound card doesn't include MIDI support. Both can convert MIDI input into WAV files that can play on any sound card. To start

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To start kmidi, click Sound & Video → More Sound & Video Applications → KMidi (or type **kmidi &** from a Terminal window). Figure 8-4 shows the KMidi window.

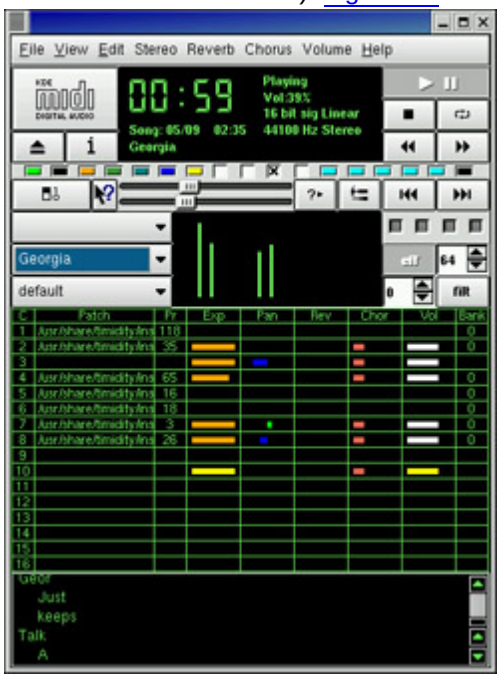

Figure 8-4: The KMidi window shows each MIDI track as it plays.

The instrument name is shown for each track (on the left), while bars (above) show the sound level for each track as it plays.

Performing audio file conversion and compression

There are many different formats for storing and compressing speech and music files. Because music files can be large, they are typically stored in a compressed format. While MP3 has been the compression format of choice, Ogg Vorbis is quickly becoming a favorite format for compressing music in the open-source community.

Tools that come with Red Hat Linux for converting and compressing audio files include:

sox — A general-purpose tool for converting audio files among a variety of formats.
oggenc — A tool for specifically converting music files to Ogg Vorbis format.

Converting audio files with Sox

If you have a sound file in one format, but you want it to be in another format, Linux offers some conversion tools you can use to convert the file. The Sox utility can translate to and from any of the audio formats listed in Table 8-1.

Table 8-1: Sound Formats Supported by Sox Utility

File Extension or Pseudonym	Description	File Extension or Pseudonym	Description
`.8svx`	8SVX Amiga musical instrument description format.	`.aiff`	Apple IIc/IIgs and SGI AIFF files. May require a separate archiver to work with these files.
`.au`	Sun Microsystems AU audio files. This is a popular format.	`.avr`	Audio Visual Research format, used on the Mac.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 8-1: Sound Formats Supported by Sox Utility

File Extension or Pseudonym	Description	File Extension or Pseudonym	Description
`.cdr`	CD-R files used to master compact disks.	`.cvs`	Continuously variable slope delta modulation, which is used for voice mail and other speech compression.
`.dat`	Text data files, which contain a text representation of sound data.	`.gsm`	Lossy Speech Compression (GSM 06.10), used to shrink audio data in voice mail and similar applications.
`.hcom`	Macintosh HCOM files.	`.maud`	Amiga format used to produce sound that is 8-bit linear, 16-bit linear, A-law, and u-law in mono or stereo.
`.ogg`	Ogg Vorbis compressed audio, which is best used for compressing music and streaming audio.	`.ossdsp`	Pseudo file, used to open the OSS /dev/dsp file and configure it to use the data type passed to Sox. Used to either play or record.
`.sf`	IRCAM sound files, used by CSound package and MixView sample editor.	`.smp`	SampleVision files from Turtle Beach, used to communicate with different MIDI samplers.
`.sunau`	Pseudo file, used to open a /dev/audio file and set it to use the data type being passed to Sox.	`.txw`	Yamaha TX-16W from a Yamaha sampling keyboard.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 8-1: Sound Formats Supported by Sox Utility

File Extension or Pseudonym	Description	File Extension or Pseudonym	Description
`.vms`	Used to compress speech audio for voice mail and similar applications.	`.voc`	Sound Blaster VOC file.
`.wav`	Microsoft WAV RIFF files. This is the native MS Windows sound format.	`.wve`	8-bit, a-law, 8 kHz sound files used with Psion Palmtop computers.
.raw	Raw files (contain no header information, so sample rate, size, and style must be given).	`.ub`, `.sb`, `.uw`, `.sw`, `.ul`, `.al`, `.lu`, `.la`, `.sl`	Raw files with set characteristics. ub is unsigned byte; sb is signed byte; uw is unsigned word; sw is signed word; and ul is ulaw.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Type **sox -h** to see the supported audio types. This also shows supported options and effects.

If you are not sure about the format of an audio file, you can add the .auto extension to the filename. This triggers Sox to guess what kind of audio format is contained in the file. The .auto extension can only be used for the input file. If Sox can figure out the content of the input file, it translates the contents to the sound type for the output file you request.

In its most basic form, you can convert one file format (such as a .wav file) to another format (such as an .au file) as follows:

$ **sox file1.wav file1.au**
To see what Sox is doing, use the `-V` option. For example:

$ **sox -V file1.wav file1.voc**

sox: Reading Wave file: Microsoft PCM format, 2 channel, 44100 samp/sec

sox: 176400 byte/sec, 4 block align, 16 bits/samp, 50266944 data bytes

sox: Input file: using sample rate 11025

size bytes, style unsigned, 1 channel

sox: Input file1.wav: comment "file1.wav"

sox: Output file1.voc: using sample rate 44100

size shorts, encoding signed (2's complement), 2 channels

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

sox: Output file: comment "file1.wav"
You can apply sound effects during the Sox conversion process. The following example shows how to change the sample rate (using the `-r` option) from 10,000 kHz to 5000 kHz:

$ **sox -r 10000 file1.wav -r 5000 file1.voc**

To reduce the noise, you can send the file through a low-pass filter. Here's an example:

$ **sox file1.voc file2.voc lowp 2200**
For more information on Sox and to get the latest download, go to the SoX — Sound eXchange — home page (`sourceforge.net/projects/sox/`).

Compressing music files with oggenc

The `oggenc` command takes music or other audio data and converts it from uncompressed formats (such as WAV, raw, or AIFF) to the compressed OGG Vorbis format. Using OGG Vorbis, audio files can be significantly reduced without a noticeable loss of sound quality. (Using default settings in `oggenc`, I reduced a 48MB WAV music file to 4MB.)
In its most basic form, you can use `oggenc` with one or more WAV or AIFF files following it. For example:

$ **oggenc *.wav**
This command would result in all files ending with `.wav` in the current directory to be converted to OGG Vorbis format. An OGG file is produced for each WAV file, with `oggenc` substituting `.ogg` for `.wav` as the file suffix for the compressed file. OGG Vorbis files can be played in many different CD players in Linux, including the xmms window (described earlier).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you want to rip music files from a CD and compress them, you can use the grip window (described later in this chapter). Grip allows you to select ogg encas the tool to do

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Cross-Reference

If you are interested in making a CD jukebox that rips, records, and compresses music CDs using oggenc and other open source software, check out the book *Linux Toys* by Christopher Negus and Chuck Wolber from Wiley Publishing.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Viewing TV and Webcams

Getting TV cards, Webcams, and other video devices to play in Linux is still a bit of an adventure. Most manufacturers of TV cards and Webcams are not losing sleep to produce Linux drivers. As a result, most of the drivers that bring video to your Linux desktop have been reverse-engineered (that is, they were created by software engineers who watched what the video device sent and received, rather than seeing the actual code that runs the device).
The first, and probably biggest, trick is to get that TV card or Webcam that is supported in Linux. Once you are getting video output from that device (typically available from `/dev/video0`) you can try out a couple of applications to begin using it. This section describes the Xawtv program for watching television and the GnomeMeeting program for video conferencing.

Watching TV with Xawtv

The Xawtv program lets you display video output, in particular television channels, on your desktop. You can change the channels, adjust volume, and fine-tune your picture. Added features let you grab still images in PBM or JPEG formats, or capture video clips in AVI.
The Xawtv will display, by default, any device producing video on the `/dev/video0` device. Therefore, you can use Xawtv to view Webcams as well as receive television channels. The following sections describe how to choose a TV capture card and use Xawtv to watch television on your desktop.

Getting a supported TV card

Video4Linux is the video interface included with Red Hat Linux. It supports a variety of TV capture cards and cameras. For TV capture cards, Video4Linux relies on the bttv driver.

To see a list of supported TV cards that you can use with Xawtv, refer to the Cards file. To view that file, you need to have the kernel-source package installed. You'll find the Cards file in the following location on your Linux system:

/usr/src/linux*/Documentation/video4linux/bttv/Cards
Also in that directory are the `CARDLIST` and `Insmod-options` files. The `CARDLIST` file lists the card number associated with each supported card. The `Insmod-options` file describes options you may want to add to the TV card driver when you load it.
Video4Linux is designed to autodetect your TV capture card and load the proper modules to activate it. So install the TV-card hardware (with the appropriate connection to your TV reception), restart Red Hat Linux, and run the `xawtv` command as described in the next section. You should be able to see video displayed on your Xawtv window.

If your card appears not to be working, here are a few things you can try:

1. To see if your TV card was properly seated in its slot and detected by Linux, type the following:
2. $ **lspci**
   This will show you a list of all valid PCI cards on your computer. If nothing shows up for the card, you probably have a hardware problem.
3. It is possible that the card is there, but the right card type is not being detected. Improper detection is most likely if you have a card for which there are several revisions, with each requiring a different driver. If you think your card is not being properly detected, find your card in the `CARDLIST` file. Then add the appropriate line to the `/etc/modules.conf` file. For example, to add a Prolink PV-BT878P, revision 9B card, add the following line to `/etc/modules.conf`:

   ```
   options  bttv  card=72
   ```

You can also add other options listed in the `Insmod-options` file. If you are still having problems getting your card to work, a mailing list is available on which you can ask questions about Video4Linux issues. The location is:

http://listman.redhat.com/mailman/listinfo/video4linux-list
One other possible reason that you don't see any video when you try to run Xawtv or other video

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

applications is that the permission of `/dev/video0` is not open. By default, that device is only open to the root user. To open the permissions on your computer so that any user can view video, type the following as root user from a Terminal window:

chmod 777 /dev/video0

Running xawtv

To start up the Xawtv viewer, simply type the following from a Terminal window on your desktop:

$ **xawtv &**

A video screen should appear in a window on the desktop. Click the left mouse button on the window to see a list of stations. Click the right mouse button to see the Settings window.

If you are in the United States, the first thing you should probably do is set the TV norm from PAL to NTSC. Then click Channel Editor and identify the names and channel numbers of each TV station. Click Update and save to make the changes permanent.
Figure 8-5 shows an example of the Xawtv window with the control window on the left and the list of stations on the right.

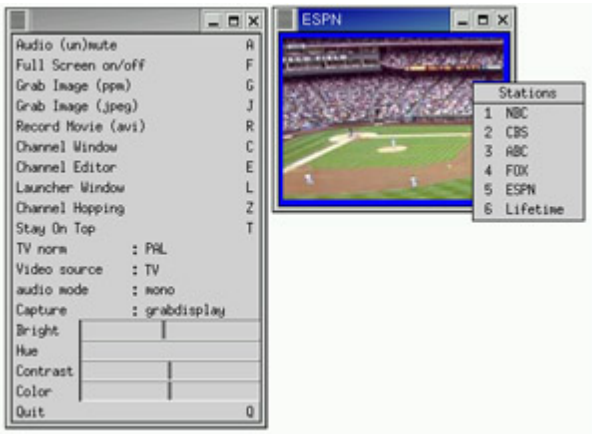

Figure 8-5: Watch TV from a TV capture card using Xawtv.

Here are a few things you can now do with your TV window:

Change channels — Use the up and down arrows to change channels.

Set up stations — Assign names to different channels so they appear in the channel list when you press the E key. You can change the name (Station ID) and Hotkey for each channel.

Grab an image — If you see an image you like on the screen, press J. A JPEG image (beginning with the word snap) appears in the directory where you started Xawtv.

Adjust picture and audio — Use function keys F6–F12 to change the picture quality. Use the plus (+) and minus (-) keys to raise or lower the volume, respectively.

Table 8-2 contains a list of the keys you can use for different functions with Xawtv.

Table 8-2: Shortcuts for TV Watching (xawtv)

Keystroke	Description
A	Audio. Mutes and unmutes audio.
F	Full Screen. Toggles between standard size and full screen.
G	Grab PPM. Captures the current image in PPM format. (Use Ctrl+G to grab the video size and not the whole window in PPM.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 8-2: Shortcuts for TV Watching (xawtv)

Keystroke	Description
H	Get Help. Brings up the Help screen that appears the first time you open xawtv.
J	Grab JPEG. Captures the current image in JPEG format. (Use Ctrl+J to grab the video size and not the whole window in JPG.)
R	Record Movie. Brings up the window that enables you to capture video clips in AVI format.
O	Settings Window. Opens the Settings window, enabling you to select options with the mouse instead of by using keystrokes.
C	Channel Window. Opens the Channel window so you can change channels.
E	Channel Editor. Opens Channel Editor so you can add stations and hotkeys.
Z	Channel Hopping. Hops to a new station every few seconds.
T	Top. Keeps display window on top.
Up/Down arrows	Channel Change. Changes channels up or down.
Left/Right arrows	Tune. Fine-tunes the image.
Page Up/Page Down	Favorites. Changes channels up and down based on a list of favorite stations saved in your configuration file.
Function keys (F5–F12)	Picture. Adjusts the brightness (F5 and F6), hue (F7 and F8), contrast (F9 and F10), and color (F11 and F12).
+/-	Volume. Adjusts the volume up or down.
Enter	Mute. Turns the volume on or off.
Q	Quit

Settings that you change in the course of using Xawtv are stored in your `$HOME/.xawtv` file. To create a listing of channels you can run the `scantv` command. It will ask about your TV and frequency tables, and then scan for channels and create a `.xawtv` file from its findings.

Video conferencing with GnomeMeeting

The GnomeMeeting window lets you communicate with other people over a network through video, audio, and typed messages. Because GnomeMeeting supports the H323 protocol, you can use it to communicate with people using other popular video-conferencing clients, such as Microsoft NetMeeting, Cu-SeeMe, and Intel VideoPhone.

To be able to send video, you need a Webcam that is supported in Linux. Although not all Webcams are supported in Linux, you still have a few dozen models to choose from. The following sections show you how to set up your Webcam and use GnomeMeeting for video conferencing.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Getting a supported Webcam

As with support for TV capture cards, Webcam support is provided through the video4linux interface. Some of the supported cameras have a parallel-port interface, although most Webcams currently supported in Linux require a USB port.
To see if your Webcam is supported, check the `/usr/src/linux*/Documentation` directory. A few parallel-port video cameras are described in the `video4linux` subdirectory; however, the bulk of the supported cameras are listed in the `usb` directory.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip After doing some research myself, I purchased a Logitech QuickCam Pro 3000. The driver for this Webcam was made for a Philips USB Webcam, but it also works for Webcams from Lo

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Supported USB cameras should be autodetected, so that when you plug them in the necessary modules are automatically loaded. Just start up the `gnomemeeting` or `xawtv` commands, and if everything is working you should see video from your Webcam on your Linux desktop. Here are a few things I typed to check that my Webcam was working properly:

```
# lsmod
pwc                43392   1
videodev            5120   2  [pwc]
usbcore            59072   1  [audio pwc usb-uhci]
```

The output from `lsmod` shows that the `pwc` driver was loaded and associated with the `videodev` module and `usbcore` module. To see information about the pwc module (which is specific to this Webcam), I typed the following `modinfo` command:

```
# modinfo –p pwc
filename:    /lib/modules/2.4.20-2.48/kernel/drivers/usb/pwc.o
description: "Philips USB webcam driver"
author:      "Nemosoft Unv. <nemosoft@smcc.demon.nl>"
license:     "GPL"
parm:        size string, description "Initial image size. One of
             sqcif, qsif, qcif, sif, cif, vga"
parm:        fps int, description "Initial frames per second. Varies
             with model, useful range 5-30"
parm:        fbufs int, description "Number of internal frame buffers
             to reserve"
parm:        mbufs int, description "Number of external (mmap()ed)
             image buffers"
parm:        trace int, description "For debugging purposes"
parm:        power_save int, description "Turn power save feature in
             camera on or off"
parm:        compression int, description "Preferred compression
             quality. Range 0 (uncompressed) to 3 (high compression)"
parm:        leds int array (min = 2, max = 2), description "LED on,off
             time in milliseconds"
parm:        dev_hint string array (min = 0, max = 10), description
             "Device node hints"
```

The output shows that pwc is a Philips USB Webcam driver. It also shows options that you can set with the pwc module in order to change the image size, frames per second, and debugging level, among other things. The information about options for your Webcam driver can be useful later if you want to tune the behavior of your Webcam or track down a problem.

Running GnomeMeeting

To start GnomeMeeting from a Terminal window, type **gnomemeeting&.** If it is not installed, you can install the gnomemeeting package from the second Red Hat installation CD (CD #2). The first time you run GnomeMeeting, the GnomeMeeting Configuration Assistant starts. The assistant lets you enter the following information:

Personal Data — Your first name, last name, e-mail address, comment, and location. You can also select whether or not you want to be listed in the GnomeMeeting ILS directory.

Connection Type — Indicate the speed of your Internet connection (56K modem, ISDN, DSL/Cable, T1/LAN or Custom).

Figure 8-6 shows the GnomeMeeting window with the history log to the right and the address book

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

behind.

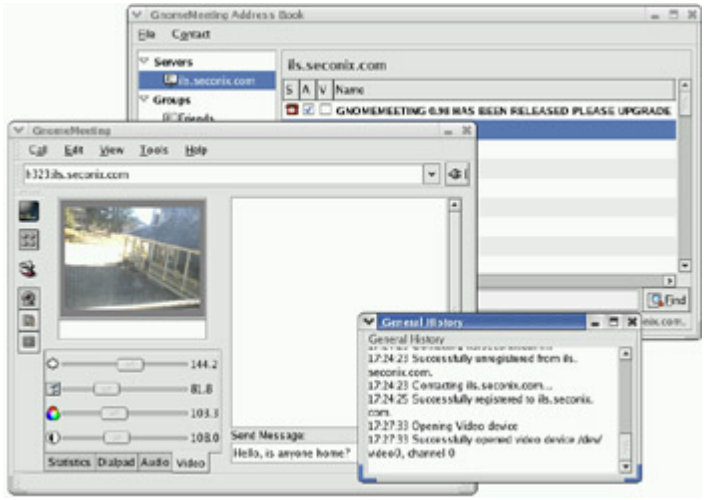

Figure 8-6: Connect to ILS servers to video-conference with GnomeMeeting.

In the GnomeMeeting window that opens you can click on the book icon to open a search window (shown to the upper right in Figure 8-6). The search window enables you to connect to an ILS directory. By typing a name into the search filter, you can search for people who might be connected to that server by first name, last name, e-mail address, or location. Select a person from the list that appears and, if he or she accepts your call, you can begin video-conferencing.

Tabs beneath the video window let you adjust your audio levels and video appearance. The History tab shows a log of your activities.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Playing Video

Although several fairly high-quality video players are available for Linux, none are included in the Red Hat Linux distribution. Legal issues surrounding the playing of encoded DVD movies in Linux might be responsible for keeping players such as the Mplayer (freshmeat.net/mplayer), Ogle (http://www.dtek.chalmers.se/groups/dvd), and Xine (xine.sourceforge.net) video players out of this distribution.

By most accounts, however, if you don't download and use the DeCCS (software for decrypting DVD movies), you can get and use these video players to play a variety of video content for personal use. The following sections provide descriptions of video applications to use with Red Hat Linux.

Cross-Reference

The mplayer video player, along with various tools for creating your own videos and TV recordings in Red Hat Linux, are available with the book *Linux Toys* by Christopher Negus and Chuck Wolber from Wiley Publishing. Those tools are described for making simple home video archive, television recorder, and streaming media projects.

Watching Video with Xine

The xine player is an excellent application for playing a variety of video and audio formats. You can get xine from xine.sourceforge.net or by downloading RPMs from freshrpms.net.
You can start the xine player by typing **xine&** from a Terminal window. Figure 8-7 shows an example of the xine video player window and controls.

Figure 8-7: Play video CDs, MP3s, Quicktime, and other video formats with xine.

Note

When you try to install xine, it will tell you if you need any additional packages. If your xine player fails to start, see the "Xine tips" section later in this chapter.

Xine supports a bunch of video and audio formats, including:

- MPEG (1, 2, and 4)
- Quicktime (see "Xine tips" if your Quicktime content won't play)
- Motion JPEG
- MPEG audio (MP3)
- AC3 and Dolby Digital audio
- DTS audio
- Ogg Vorbis audio

Xine can understand different file formats that represent a combination of audio and video. These include .mpg (MPEG program streams), .ts (MPEG transport streams), .mpv (raw MPEG audio/video streams), .avi (MS AVI format), and .asf (Advanced Streaming format). While xine can play Video CDs and DVDs, it can't play encrypted DVDs or Video-on-CD hybrid format as xine is delivered (because of the legal issues mentioned earlier related to decrypting DVDs).

Using xine

With xine started, right-click in the xine window to see the controls. The quickest way to play video is to click one of the following buttons, then press the Play button (right arrow or play, depending on the skin

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

you are using):

VCD (looks for a video CD)
DVD (looks for a DVD in `/dev/dvd`)
CDA (looks a music CD in `/dev/cdaudio`)

Next, you can use the Pause/resume, Stop, Play, Fast motion, Slow motion, or Eject buttons to work with video. You can also use the Previous and Next buttons to step to different tracks. The controls are very similar to what you would expect on a physical CD or DVD player.

To select individual files, or to put together your own list of content to play, you can use the Playlist feature.

Creating playlists with xine

Click the Playlist button on the left side of the xine control window. A Playlist Editor appears, showing the files on your current playlist. You can add and delete content from this list, then save the list to call on later. Here's how you use the xine Playlist Editor:

CDA, DVD, or VCD — Click any of the buttons that represent a particular CD or DVD. All content from that CD or DVD is added to the playlist.
Add — Click the Add button to see the MRL Browser window. From that window, click File to choose a file from your Linux file system to add to the list. Click Select to add that file to the Playlist Editor.
Move up/ Move down — Use the Move up selected MRL and Move down selected MRL buttons to move up and down the playlist.
Delete — Click the Delete Selected MRL button to remove the current selection.
Delete all — Click the Delete All Entries button to clear the whole playlist.
Save — Click the Save button to save the playlist to your home directory (`$HOME/.xine/playlist`).
Load — To read in the playlist you saved, click the Load button.

The xine content is identified as media resource locators (MRLs). Each MRL is identified as a file, DVD, or VCD. Files are in the regular file path (`/path/file`) or preceded by a file `:/`, `fifo:/`, or `stdin:/`. DVDs and VCD are preceded by `dvd` and `vcd`, respectively (for example, `vcd://01`).

To play your playlist, click the Play button (arrow key) on the Playlist Editor.

Xine tips

Getting video and audio to work properly can sometimes be a tricky business. Here are a few quick tips if you are having trouble getting xine to work properly (or at all):

Xine won't start — To work best, xine needs an X driver that supports OpenGL. If there is no OpenGL support for your video card in X, xine will shut down immediately when it tries to open the default Xv driver. If this happens to you, try starting xine with the X11 video driver (which is slower, but should work) as follows:
$ **xine -VXSHM**
Xine playback is choppy — If playback of files from you hard disk is choppy, there are a couple of settings you can check: 32-bit IO and DMA. (If these two features are supported by your hard disk, they will generally improve hard disk performance.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

Improper disk settings can result in destroyed data on your hard disk. Do this procedure at your own risk. This procedure is only for IDE hard drives (NO SCSI)! Also, be sure to

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

First, test the speed of hard disk reads. To test the first IDE drive (`/dev/hda`), type:

hdparm -t /dev/hda
Timing buffered disk reads: 64 MB in 19.31 seconds = 3.31 MB/sec

To see your current DMA and IO settings, as root user type:
hdparm -c -d /dev/hda
/dev/hda:
 I/O support = 0 (default 16-bit)
 using_dma   = 0 (off)

This shows that both 32-bit IO and DMA are off. To turn them on, type:
hdparm -c 1 -d 1 /dev/hda
/dev/hda:
 I/O support = 1 (32-bit)
 using_dma   = 1 (on)

With both settings on, test the disk again:

hdparm -t /dev/hda
Timing buffered disk reads: 64 MB in 2.2 seconds = 28.83 MB/sec

As you can see from this example, buffered disk reads of 64 MB went from 19.31 seconds to 2.2 seconds after changing the parameters described. Playback should be much better now.

Xine won't play particular media — Messages such as "no input plug-in" mean that either the file format you are trying to play is not supported or it requires an additional plug-in (as is the case with playing DVDs). If the message is "maybe xyx is a broken file," the file may be a proprietary version of an otherwise supported format. For example, I had a Quicktime video fail that required an SVQ3 codec (which is currently not supported under Linux), though other Quicktime files will play fine.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The Cross Over Plugin (described in Chapter 9) can be used to play a variety of content, including the version of Quicktime just mentioned.

Using RealPlayer

A tremendous amount of content is available on the Internet in the RealMedia and RealAudio formats. You can see and hear video clips of popular musicians and comics. You can view live events, such as conferences, news stories, and concerts. You can also listen to your favorite radio stations when you are out of town.

To play RealMedia and RealAudio content you need, as you may have guessed, the RealPlayer. Real Networks (`www.real.com`) is a leader in streaming media on the Internet. More than 50 million unique users have registered with Real Networks and their Web site, downloading more than 175,000 files per day. And that's not even the good news. The good news is that a RealPlayer is available to run in Red

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Hat Linux.
The RealPlayer for Linux is available via the Linux area of download.com or tucows.com. Or, try this site: http://proforma.real.com/real/player/unix/unix.html. This player is not supported by Real Networks directly.

The instructions for configuring RealPlayer are delivered in HTML format, so you can read it in Netscape or some other Web browser. If any patches or workarounds are required, you can find them in the Real Networks Knowledge Base. To get there, click Support (from most Real Networks pages), then click Knowledge Base. When there, query for the word Linux to find any problem reports and fixes.

When you install the RealPlayer, you are asked if you want to configure it to be used as a Netscape plug-in (which I recommend you do so you can play real content in Mozilla). After that, when you open any Real content in your browser, the RealPlayer opens to handle it. Alternatively, you can start the RealPlayer from a Terminal window on your desktop by typing the following:

$ **realplay &**
Real Networks has gone to a subscription model for its content. You sign up and pay a monthly fee to get RealPlayer content. To see what is available, and to decide if it is worth signing up, I suggest starting at the RealGuide site (realguide.real.com). Besides describing the Real Networks content that is available, there are a few clips at this site you can try out.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using a Digital Camera with gtkam and gphoto2

With the gtkam window, you can download and work with images from digital cameras. The gtkam window is front end to gPhoto2, which provides support for dozens of digital cameras in Linux. The gtkam window works by attaching a supported digital camera to a serial or USB port on your computer. You can view thumbnails of the digital images from the camera, view full-size images, and download the ones you select from the camera to your hard disk.

Note If you have a camera that saves images to a floppy disk, just insert that disk into your disk drive and the contents of the disk should open automatically on your desktop.

Check the gPhoto2 Web site (`http://www.gphoto.org/proj/libgphoto2/support.php`) for information on supported cameras as well as other topics related to gPhoto. Here is a list of currently supported digital cameras.

Agfa ePhoto — Supported models: 307, 708, 708C, 1280, and CL18.
Apple QuickTake — Supported model: 200.
Cannon Digital — Suported models: IXUS, IXUS 300, and IXUS v.
Cannon PowerShot — Supported models: A5, A5 Zoom, A10, A20, A50, A70, G1, G2, Pro70, Pro90 IS, S10, S100, S110, S300, and S20.
Casio QV — Supported models: 10, 10A, 11, 30, 70, 100, 200, 700, and 5000SX.
Chinon — Supported model: ES-1000.
Epson PhotoPC — Supported models: 500, 550, 600, 700, and 800.
Fuji — Supported models: DS-7, DX-5, DX-10, MX-500, MX-600, MX-700, MX-1200, and MX 2700.
Hewlett-Packard PhotoSmart — Supported models: 618, C20, C30, and C2000.
Kodak DC — Supported models: 120, 20, 25, 200+, 210, 210+ Zoom, 215 Zoom, 220 +, 240, 260, 265, 280, 290, 3200, 3400, and 5000.
Konica — Supported models: QM100, QM100V, Q-EZ, Q-M100, Q-M100V, Q-M200.
Leica — Supported model: Digilux Zoom.
Minolta — Supported model: Dimage V.
Mustek — Supported model: MDC 800 v2.
Nicon CoolPix — Supported models: 100, 300, 600, 700, 800, 880, 900, 900S, 910, 950, 950S, and 990.
Olympus — Supported models: D-100Z, D-200L, D-220L, D-300L, D-320L, D-330R, D-340L, D-340R, D-400L Zoom, D-450Z, D-500L, D-600L, D-620L, C-400L, C-410L, C-800L, C-820L, C-830L, C-840L, C-900 Zoom, C-900L Zoom, C-1000L, C-1400L, and C-2000Z.
Panasonic — Supported models: Coolshot KXL-600A and KXL-601A, and NV-DCF5E, and DC1000 and DC1580.
Philips — Supported models: ESP60 and ESP 80.
Polaroid — Supported models: PDC 640 and DC700.
Ricoh RDC — Supported models: 300, 300Z, 4200, 4300, and 5000.
Samsung — Supported models: Kenox SSC-350N and Digimax 800K.
Sanyo VPC — Supported models: G210, G200, G200EX, G250, X300, and X350.
Sony — Supported models: DSC-F1, DSC-F55, DSC-F505, Memory Stick Adapter, and MSAC-SR1 and DCR-PC100.
Toshiba — Supported model: PDR-M1.

Downloading digital photos with gtkam

The following procedure describes how to download images from your digital camera.

1. Using a cable provided with your digital camera, connect your camera to the USB or COM port on your computer. (I had better luck with the USB port.)
2. Set your camera to be in Send and Receive mode.
3. From the main Red Hat desktop menu, choose Graphics → Digital Camera Tool. The gtkam window appears.
4. Click Camera → Add Camera. The Select Camera window appears.
5. Click the down arrow next to the Model box, select your camera, and click Detect.
6. Click Apply, then OK. Your camera model should be listed in the gtkam window.
7. To begin downloading images from your digital camera, click the camera name that

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

appears in the left column, and then select the folder containing the images from that camera. After the images download (which can take a while), thumbnails appear in the main gtkam window, as shown in Figure 8-8.

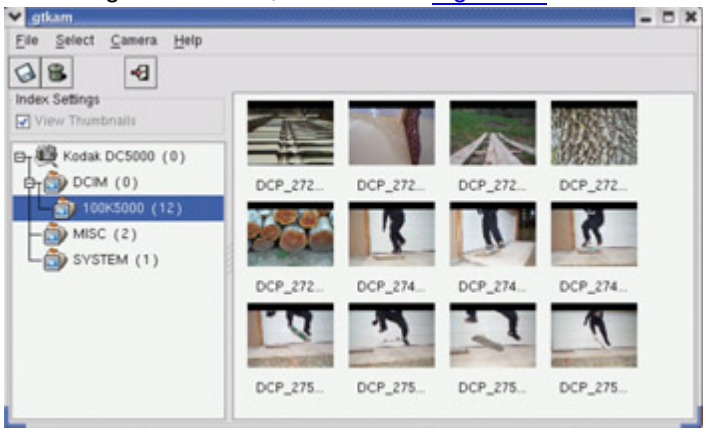

Figure 8-8: Download images from digital cameras from the gtkam window.

8. Select the images you are interested in, and click the Save Selected Photos button to save the selected images. The Save Photos window that appears lets you choose a directory to save them to. You can rename the images or just use the names assigned by the camera.
9. Choose images you want to delete, and click the Delete button to delete them.

Using your camera as a storage device

Instead of using gtkam or gphoto to download pictures to Linux, some digital cameras let you treat them like a storage device to manage pictures. By mounting a digital camera as a USB mass storage device, you can view, copy, delete, and move the pictures on your camera as you would files on a hard disk or CD (just at a lower speed).

The following is a partial list of digital cameras that can be used as a USB storage device:

Casio — Supported models: QV-2400UX, QV-2x00, QV-3x00, QV-4000 and QV-8000.
Fuji — FinePix 1300, 1400Zoom, 2300Zoom, 2400Zoom, 2800Zoom, 4200Z, 4500, 4700 Zoom, 4900 Zoom, 6800 Zoom, A101, A201, and S1 Pro.
HP — PhotoSmart 315, 318xi, 618, and C912,
Konica — KD200Z, KD400Z, and Revio KD300Z.
Kyocera — Finecam s3.
Leica — Digilux 4.3.
Minolta — Dimage 5, Dimage 7, and Dimage X.
Nikon — CoolPix 2500, 885, 5000, 775, and 995.
Olympus — Brio Zoom D-15, C-100, C-200Z, C-2040, C-220Z, C-2Z, C-3020Z, C-3040Z, C-4040Zoom, C-700, C-700UZ, C-860L, D-510, D-520Z, E-10, and E-20.
Pentax — EI2000, Optio 330, and Optio 430.
Sony — DSC-F505, DSC-F505V, DSC-F707, DSC-P1, DSC-P20, DSC-P5, DSC-P71, DSC-S30, DSC-S70, DSC-S75, DSC-S85, MVC-CD300, and MVC-FD92
Vivitar — Vivicam 3550
Yashica — Finecam s3

To Linux, the USB mass storage camera appears as a SCSI drive containing a VFAT file system with image files on it. Here's a procedure for using your digital camera as a USB storage device:

1. Using a cable provided with your digital camera, connect your camera to a USB port on your computer and turn it on so it is ready to send and receive data.
2. Boot your computer to Red Hat Linux.
3. Open the `/etc/fstab` file as root user and see if an entry was created for your digital camera. If you have no other SCSI devices on your computer, the camera will probably be detected as `/dev/sda1` device. Here is what the entry might look like:

   /dev/sda1 /mnt/camera auto defaults, user,noauto 0 0
   If no such entry appears, create the entry, possibly using `vfat` instead of auto. Create the mount point directory (as root user, type `mkdir /mnt/camera`).
4. As root user, type the mount command to mount the camera: **mount /mnt/camera**.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

5. Open the `/mnt/camera` as you would any other directory from the shell or from a file manager. Copy, delete, move, and rename files as you would any files on your hard disk.
6. When you are done, be sure to unmount the camera as follows (as root user from a Terminal window).
7. # **umount /mnt/camera**

Caution If you unplug your camera without unmounting the file system, it could damage the files on your camera.

You can follow the procedure above to use other USB mass storage devices (CD drives, keychains, and so on) in Linux. Use different mount directories (such as `/mnt/keychain`) and check which SCSI device is being assigned to the USB storage device.

To see if your USB storage device can be seen by Linux, you can check the `/var/log/dmesg` file or run the `usbview` command. Either of those places will tell you if the device is being detected properly by Linux.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Recording Music CDs

Writable CD-ROM drives are fast becoming a standard device on computers. Where once you had to settle for a floppy disk (1.44MB) or a Zip disk (100MB) to store personal data, a CD-ROM burner lets you store more than 600MB of data in a format that can be exchanged with most computers. On top of that, you can create CD music disks!
Both graphical and command-line tools exist for creating CDs in Red Hat Linux. The `cdrecord` command lets you create audio and data CDs from the command line. Using cdrecord you can write to CD-Recordable (CD-R) and CD-Rewritable (CD-RW) drives. The `cdrecord` command is described in the next section of this chapter.

Configuring an IDE CD-ROM for recording

If you are using a SCSI drive, you are probably ready to begin recording; CD recording tools in Linux should recognize and be able to use your drive automatically. If you are using an IDE-based drive, however, you may have some work to do.

In most cases, your IDE writable CD drive should be recognized and properly configured when you install Red Hat Linux. The following boot option should automatically appear as an option at the end of the kernel line on your bootloader screen:

hdc=ide-scsi

This shows that IDE disk device "c"(hdc) is set up to do SCSI emulation. If this is working properly, you should be able to type the following to see if the CD drive appears on the list of accessible SCSI devices.

$ **cdrecord --scanbus**

If your IDE writable CD drive is not working, check that the following basic steps are done: 1) have SCSI emulation built into the kernel; 2) remove the IDE-CD module; and 3) add the IDE-SCSI module. By default, SCSI emulation should be built into your kernel, so unless you remove it, there's nothing you have to do there. Here's how you do the other two steps:

1. To remove the ide-cd module, as root user add the following line to the `/etc/modules.conf` file (using any text editor).
   options ide-cd ignore=hdc
2. In this example, the writable CD-ROM device is located at `/dev/hdc`. Replace `hdc` with the location of your CD device. If you are not sure, type **dmesg | grep ^hd** to see a list of IDE disks on your computer.
3. Restart your computer.
4. To add the ide-scsi module, type the following command:
5. # **modprobe -a ide-scsi**
6. To make sure Linux can recognize your CD writer, type the following command:
7. # **cdrecord -scanbus**
8. Cdrecord 1.10 (i686-pc-linux-gnu) Copyright (C) 1995-2001
9. Jörg Schilling
10. Linux sg driver version: 3.1.24
11. Using libscg version 'schily-0.5'
12. scsibus0:
13. 0,0,0 0) 'IDE-CD' 'R/RW 4x4x24 "1.04' Removable CD-ROM
14. 0,1,0 1) 'SONY "CD-ROM CDU701 "1.0f' Removable CD-ROM
15. 0,2,0 2) *
16. 0,3,0 3) *
17. 0,4,0 4) *
18. 0,5,0 5) *
19. 0,6,0 6) *
    0,7,0 7) *

    Here you can see that `cdrecord` has detected two CD drives on SCSI bus 0. A writable CD drive is at 0,0,0 and a read-only CD is located at 0,1,0. You need to remember these numbers later when you use `cdrecord` to create a CD. Until your CD burners show up on this list, `cdrecord` cannot write to them.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

20. After you have added your CD writer as a SCSI device, it may no longer be available to the CD players described in this chapter. By default, those players play the `/dev/cdrom` device, which in this case points to an IDE drive. To fix that, type the following as root user:
21. # **cd /dev**
22. # **rm cdrom**
23. # **ln -s scd0 cdrom**
24. If you have two CD drives, repeat this step, replacing cdrom with cdrom1, and scd0 with scd1.

Creating an audio CD with cdrecord

You can use the `cdrecord` command to create either data or music CDs. You can create a data CD by setting up a separate file system and copying the whole image of that file system to CD. Creating an audio CD consists of selecting the audio tracks you want to copy and copying them all at once to the CD.

This section focuses on using `cdrecord` to create audio CDs. The `cdrecord` command can use audio files in `.au`, `.wav`, or `.cdr` format, automatically translating them when necessary. If you have audio files in other formats, you can convert them to one of the supported formats by using the `sox` command (described previously in this chapter).

Cross-Reference See Chapter 13 for how to use cdrecord to create data CDs.

One way to create an audio CD is to copy the music tracks you want to a directory; then copy them to the writable CD. To extract the tracks, you can use the `cdda2wav` command. Then you write them to CD by using the `cdrecord` command. Here's an example:

1. Create a directory to hold the audio files, and change to that directory. Make sure the directory can hold up to 660MB of data (or less if you are burning fewer songs). For example:
2. # **mkdir /tmp/cd**
3. # **cd /tmp/cd**
4. Insert the music CD into your CD-ROM drive. (If a CD player opens on the desktop, close it.)
5. Extract the music tracks you want by using the `cdda2wav` command. For example:
6. # **cdda2wav -D /dev/cdrom -B**

   This example reads all of the music tracks from the CD-ROM drive located at SCSI device number 0,0,0. The `-B` option says to output each track to a separate file. By default, the `cdda2wav` command outputs the files to WAV audio format.

   Instead of extracting all songs, you can choose a single track or a range of tracks to extract. For example, to extract tracks 3 through 5, add the `-t3+5` option. To extract just track 9, add `-t9+9`. To extract track 7 through the end of the CD, add `-t7`.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you have a low quality CD drive or an imperfect CD, cda2wav might not be the best ripping tool. Instead of cda2wav, you could use cdparanoia-B to extract songs from the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

7. When `cdda2wav` is done, insert a blank CD into your writable CD drive.
8. Use the `cdrecord` command to write the music tracks to the CD. For example:
9. # **cdrecord -v dev=0,0,0 -audio *.wav**
   The options to `cdrecord` tell the command to create an audio CD (`-audio`) on the writable CD device located on 0,0,0. The `cdrecord` command writes all files from the current directory that end in `.wav`. The `-v` option causes verbose output.
   If you want to change the order of the tracks, you can type their names in the order you want them written (instead of using `*.wav`). If your CD writer supports higher speeds, you can use the speed option to double (`speed=2`) or to quadruple (`speed=4`) the writing speed.

Once you have created the music CD, indicate the contents of the CD on the front of the CD. The CD should now be ready to play on any standard music CD player.

Ripping CDs with Grip

The Grip window provides a more graphical method of copying music from CDs to hard disk. You can then play the songs directly from your hard disk or burn them back on to a blank CD. Besides just ripping music, you can also compress each song as you extract it from the CD.
You can open Grip from the Red Hat menu by selecting Sound & Video → More Sound & Video Applications → Grip (or by typing **`grip`** from a Terminal window). Figure 8-9 shows an example of the Grip window.

Figure 8-9: Rip and play songs from the Grip window.

To rip audio tracks from a CD with `grip`, do the following:

1. With the Grip window open, insert a music CD into your CD drive. If you have an active connection to the Internet and the CD is known to the CD database, then title, artist, and track information appear for the CD.
2. Right-click on each track that you want to rip (that is, copy to your hard disk). A check mark appears under the Rip column for the song.
3. Click the Config tab at the top of the page, then select Encode.
4. You can select the type of encoder used to compress the music by clicking the Encoder box and selecting an encoder (by default, `oggenc` compresses files in Ogg Vorbis).
5. Click the Rip tab at the top of the page. The Rip tab appears.
6. Click one of the following:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Rip+Encode — This rips the selected songs and (if you did the step above) compresses them in Ogg Vorbis format. You need an Ogg Vorbis player to play the songs after they have been ripped in this format (there are many Ogg Vorbis players for Linux).
Rip only — This rips the selected songs in WAV format. You can use a standard CD player to play these songs. (When I tried this, the same song ripped in WAV was twelve times larger than the Ogg Vorbis file.)

Songs are copied to the hard disk in the format you selected. By default, the files are copied into a subdirectory of `$HOME/ogg` (such as `/home/jake/ogg`). The subdirectory is named for the artist and CD. For example, if the user jake were ripping the song called High Life by the artist Mumbo, the directory containing ripped songs would be `/home/jake/ogg/mumbo/high_life`. Each song file would be named for the song (for example, `fly_fly_fly.wav`).

At this point you can play any of the files using a player that can play WAV or Ogg files, such as xmms. Or you can copy the files back to the CD using `cdrecord`. Because the filenames are the song names, they don't appear in the same order as they appear on the CD. So if you want to copy them back to a writable CD in the same order where they originally appeared, you may have to type each filename on the `cdrecord` command line. For example:

cdrecord -v dev=0,0,0 -audio fly_fly.wav big_news.wav about_time.wav

The Grip window can also be used to play CDs. Use the buttons on the bottom of the display to play/pause, skip ahead/back, stop, and eject the CD. The Toggle track display button lets you shrink the size of the display so it doesn't take up much space on the desktop. Click Toggle disc editor to see and change title, artist, and track information.

Creating CD labels with cdlabelgen

The `cdlabelgen` command can be used to create tray cards and front cards to fit in CD jewel cases. You gather information about the CD and `cdlabelgen` produces a PostScript output file that you can send to the printer. The cdlabelgen package also comes with graphics (in `/usr/share/cdlabelgen`) that you can incorporate into your labels.
Here is an example of a `cdlabelgen` command line that you can use to generate a CD label file in PostScript format. (Type it all on one line or use backslashes, as shown here, to put it on multiple lines.)

```
cdlabelgen -c "Grunge is Gone" -s "Yep HipHop" \
-i "If You Feed Me%Sockin Years%City Road%Platinum and Copper%Fly Fly \
Fly%Best Man Spins%What A Headache%Stayin Put Feelin%Dreams Do Go \
Blue%Us%Mildest Schemes" -o yep.ps
```

In this example, the title of the CD is indicated by `-c "Grunge is Gone"` and the artist by the `-s "Yep HipHop"` option. The tracks are entered after the `-i` option, with each line separated by a `%` sign. The output file is sent to the file `yep.ps` with the `-o` option. To view and print the results, you can use the `gv` command as follows:

$ **gv yep.ps**

The results of this example are shown in Figure 8-10.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

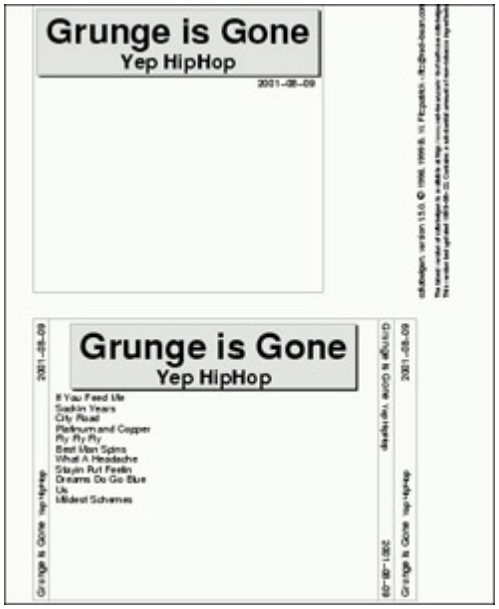

Figure 8-10: Generate CD jewel case labels with cdlabelgen and print them with gv.

You will probably want to edit the command and re-run `gv` a few times to get the CD label correct. When you are ready to print the label (assuming you have a printer configured for your computer), click Print All to print the label.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version

Please register to remove this banner.

http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

Getting your Red Hat Linux system set up for sound and video can take some doing, but once it's done you can play most audio and video content that is available today. This chapter took you through the steps of setting up and troubleshooting your sound card. When that was done, you learned how to find software to play music through that card.

Live video from TV cards and Webcams was discussed in the sections on Xawtv and GnomeMeeting, respectively. The xine player is described for playing a variety of video formats. To download images from a digital camera, the gtkam window was described. If your computer has a CD burner, use descriptions in this chapter to create your own music CDs and CD labels.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 9: Tools for Using the Internet and the Web

With your Red Hat Linux system connected to the Internet, you can take advantage of dozens of tools for browsing the Web, downloading files, getting e-mail, and participating in newsgroups. In most cases, you have several choices of GUI and command-line applications for using Internet services from your Linux desktop or shell.

This chapter describes some of the most popular tools available with Red Hat Linux for working with the Internet. These descriptions include Web browsers, e-mail readers, newsreaders, instant messaging clients, and commands for login and remote execution.

Overview of Internet Tools

The most important client Internet program these days is the Web browser. Red Hat Linux includes the Mozilla software package, which includes a Web browser along with other Web client software for reading mail, participating in newsgroups, and creating Web pages (to name a few). Other Web browsers, some of which incorporate Mozilla features, also come with Red Hat Linux. These include:

Epiphany — This Web browser is integrated with the GNOME desktop, allowing you to take advantage of GNOME themes, drag-and-drop, and translation features. On the inside, Epiphany relies on Mozilla's rendering engine.

Konqueror — Although Konqueror is the file manager for the KDE desktop, it can also display Web content. Using Konqueror, you can easily go back and forth between Web sites and local files and folders.

Running a close second to Web browsers is the e-mail reader (referred to in network standards terms as a Mail User Agent, or MUA). Ximian Evolution is the recommended e-mail client for Red Hat Linux. Another option is the Mozilla integrated e-mail client. Mail programs that have been around in Linux and other UNIX systems since the time when most mail was plain text include `mutt`, `pine` and `mail`.

You can choose from thousands of newsgroups to participate in discussions on the Internet. Red Hat Linux has several newsreaders available. Again, Mozilla includes an application for participating in newsgroups. Also, the `pan` and `slrn` newsreaders are available.

Besides browsing, e-mail, and news, there are many ways of communicating with other computers and users on the Internet. Older UNIX commands such as `rlogin`, `rsh`, and `rcp` are supported in Linux to do remote login, run remote commands, and copy files remotely. In recent years, OpenSSH commands (`ssh`, `scp`, and `sftp`) have become preferable to the "r" commands because they offer greater security.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Browsing the Web

Although the Internet has been around since the 1960s, the Web is a relatively new technology (it was created in 1985). The Web places an additional framework over Internet addresses that were once limited to host names and domain names. Before the Web, finding resources on the Internet was difficult. However, the Web now provides several features that make it much easier to access these resources:

Uniform Resource Locators (URLs) — URLs identify the location of resources on the Web. Besides identifying the domain and host on which a resource resides, they can also identify the type of content and the specific location of the content.

Hypertext Markup Language (HTML) Web pages — When people talk about a Web page, they are generally referring to information that is presented in HTML format. HTML changed the Internet from a purely plain-text-based resource to one that could present graphics and font changes. An HTML page can also contain hypertext links. Links are the threads that join together the Web, enabling someone viewing a Web page to be immediately transported to another Web page (or other content) by simply selecting a linked text string or image in the page.

The primary tool for displaying HTML Web pages is the Web browser. Mozilla is the most popular Web browser for Red Hat Linux. It displays HTML (Web pages) as well as other types of Web content. Now, even file managers, made for displaying local files and folders, have been extended to be able to display Web content (see the description of Konqueror file manager in Chapter 3).

This section contains general information about the Web and some specific hints for using Mozilla to browse the Web from your Red Hat Linux system.

Uniform Resource Locators

To visit a site on the Internet, you either type a URL into the Location box on your browser or click on a link (either on a Web page or from a menu or button on the browser). Although URLs are commonplace these days — you can find them on everything from business cards to cereal boxes — you may not know how URLs are constructed. The URL form is as follows:

protocol://host-domain/path

The protocol identifies the kind of content that you are requesting. By far, the most common protocol you come across is Hypertext Transfer Protocol (HTTP). HTTP is the protocol used to request Web pages. In addition to HTTP, however, there are other protocols that might appear at the beginning of a Web address. Instead of showing you a Web page, these other types of protocols may display different kinds of information in your browser, or open a completely different application for working with the content.

Table 9-1 lists some of the protocols that can appear in a Web URL.

Table 9-1: Protocols in Web URLs

Protocol Name	Description
`http`	Hypertext Transfer Protocol. Used to identify HTML Web pages and related content.
`file`	Identifies a file on a specific host. Most often used to display a file from your local computer.
`ftp`	File Transfer Protocol. Identifies a location where there are file archives from which you might want to download files.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 9-1: Protocols in Web URLs

Protocol Name	Description
`gopher`	Gopher Protocol. Provides databases of text-based documents that are distributed across the Internet. (Gopher is nearly obsolete.)
`mailto`	Electronic Mail Address. Identifies an e-mail address, such as `mailto://joe@example.com`. (Usually opens a mail composer.)
`news`	USENET newsgroup. Identifies a newsgroup, such as `news://news.myisp.com/comp.os.linux.networking`. If you type this address into Mozilla, a window appears with the newsgroup displayed from the news server you identified.
`nntp`	USENET news using `nntp` protocol.
`telnet`	Log in to a remote computer and begin an interactive session. An example of a telnet address is `telnet://localhost`. (Replace `localhost` with any host or IP address that allows you to log in.)
`wais`	Wide Area Information Server protocol. A WAIS address might look like the following: `wais://handsonhistory.com/waisdb`. (Like gopher, WAIS databases are nearly obsolete.)

The first part of a URL is the protocol. You don't always have to type the protocol. Most browsers are good at guessing the content you are looking for (mostly it guesses HTTP). If the address you type starts with `www`, it assumes HTTP; if it starts with `ftp`, it assumes FTP.

The second part of a URL takes you to the computer that is hosting the Web content. By convention, Web servers begin with `www` (or sometimes `home`). However, if you type the correct protocol (usually `http`), you will be directed to the right service at the host computer. The next piece of this name is just the host.domain style of Internet address that is always used with the Internet (such as `redhat.com`, `linuxtoys.net`, or `whitehouse.gov`). An optional port number can be tacked on to the host.domain name. For example, to request the port used for HTTP services (port 80) from the host called `www.linuxtoys.net`, you can type `http://www.linuxtoys.net:80`.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip You can identify a specific port number to request the service attached to the port on the computer you request. A port number is a lot like a telephone extension in a big

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The third part of a URL identifies the location of the content on the host computer. Sections in a Web page can be identified with a pound sign (#) and an identifier following the Web page location. For example, the craft section of the `bsched.htm` page at `handsonhistory.com` would appear as `http://www.handsonhistory.com/bsched.htm#craft`. The filename extension (such as `.htm` or `.html`) further identifies the content type.

Web pages

If you look at the HTML source code that produces Web pages, you see that it consists of a combination of information and markup tags, all of which are in plain-text format. The idea was to have Web pages be very portable and flexible. You can create a Web page with `vi`, `emacs`, Notepad, or any text editor on any computing platform. Or simplified front-end programs can be used to provide WYSIWYG (What You See Is What You Get) interfaces that let you see what you are creating as you go.

HTML tags are set apart by right and left angle brackets. Tags come in pairs, with a beginning tag, the information, and then an ending tag. The beginning tag contains the tag name, while an ending tag contains a forward slash (/) and the tag name. Here is a minimal HTML page:

<HTML>

<HEAD>

<TITLE>Greetings from Washington</TITLE>

</HEAD>

<BODY>

Here we are in beautiful Tacoma.

</BODY>

</HTML>

You can see that the document begins and ends with HTML tags (`<HTML>` and `</HTML>`). The beginning part of the Web page is contained within the `HEAD` tags. The body of the page is contained within the `BODY` tags. The title of the page is set apart by `TITLE` tags.

Between the beginning and ending `BODY` tags, you can add all kinds of stuff. You can have different types of bulleted or numbered lists. You can have headings, images, and text. More complex pages can include forms, dynamic HTML (which changes the content as you move or select items), or special data. Figure 9-1 is an example of a Web page as it appears in Mozilla.

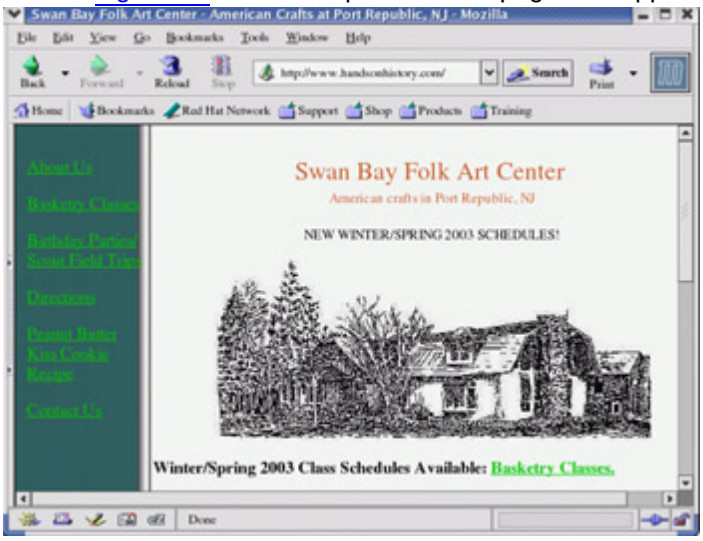

Figure 9-1: Many Web pages contain text, images, headings, and links.

Some of the HTML code that was used to create the Web page shown in Figure 9-1 is shown here. The title of the Web page appears between two `TITLE` tags.

<TITLE>Swan Bay Folk Art Center - American Crafts

in Port Republic, NJ</TITLE>

The following code is used to create a link that opens a new mail message window:

A HREF="mailto:webmaster@handsonhistory.com">Contact Us</A>;

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The text `Contact Us` is a link to an e-mail address. When someone clicks on that link, a new message window appears, allowing that person to send e-mail to that address.

<FONT FACE="Copperplate Gothic Bold"><A HREF=" bsched.htm">

Basketry Classes</A></FONT>

A special font face was used for the words `Basketry Classes`, with the words pointing to a link to another HTML page. If someone were to click on those words, the `bsched.htm` page would appear.

Web browsing with Mozilla

During the early 1990s, Netscape Navigator was the most popular Web browser. When it became apparent that Netscape was losing its lead to MS Internet Explorer, the source code for Netscape was released to the world as open-source code.

Mozilla.org was formed to coordinate the development of the new browser from that code. The result was the Mozilla browser that is now available with Red Hat Linux, as well as with many other computing platforms. Recently, the Mozilla Foundation was formed (`www.mozillafoundation.org`), to gain nonprofit status for the Mozilla project going forward.

At the center of Mozilla, of course, is the Navigator Web browser. Mozilla also includes the following features:

Mail and Newsgroups — A full-featured program for sending, receiving, and managing e-mail, as well as for using newsgroups. (The mozilla-mail RPM must be installed.)

IRC Chat — An Internet Relay Chat (IRC) window, called ChatZilla, for participating in online, typed conversations. (The mozilla-chat package must be installed.)

Composer — A Web page (HTML) composer application.

Address Book — An application to manage names, addresses, and telephone numbers. (This feature is also part of the mozilla-mail RPM package.)

You can click the Mozilla Web Browser icon directly on the desktop panel to start Mozilla. Or you can open Mozilla from the desktop menu (beginning from the red hat) by selecting Internet → Mozilla Web Browser. Figure 9-2 shows an example of the Mozilla browser as it displays the Mozilla home page (`www.mozilla.org`).

Figure 9-2: Mozilla is the open-source Web browser created from Netscape source code.

Mozilla has all the basic features you need in a Web browser and a few special features. The following sections describe how to get the most out of your Mozilla Navigator Web browser.

Setting up Mozilla Navigator

There are many things you can do to configure Mozilla to run like a champ. The following sections describe some ways to customize your browsing experience in Mozilla Navigator.

Setting Navigator preferences

You can set your Mozilla Navigator preferences in the Preferences window. To open Mozilla preferences, click Edit → Preferences. The Preferences window appears, as shown in Figure 9-3.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

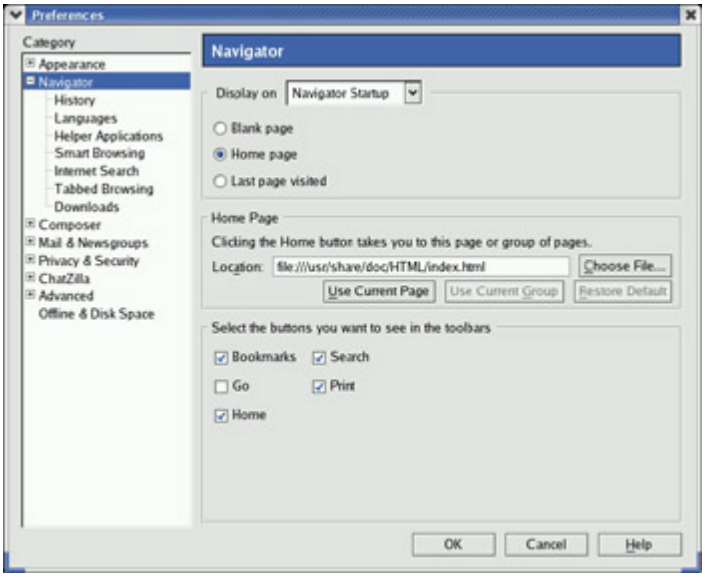

Figure 9-3: Change settings for navigating the Web from Mozilla's Preferences window

Here are some Navigator preferences that you might want to change:

Navigator — Besides selecting the location to use as your home page, you can choose which buttons you see on the toolbar.

History — Choose how long to store addresses of pages you visit. You can also select to clear your history and/or the sites you have typed in your location bar.

Languages — Set a list of preferences for the particular language a Web page should be displayed in, if the page is available in several languages.

Helper Applications — Set up applications to use to handle different types of data that may be encountered while browsing the Web.

Smart Browsing — Choose to do keyword searches when you type partial addresses in the Location box. Smart browsing is off by default, so typing `netscape` would cause Mozilla to look for `www.netscape.com` instead of searching for `netscape` links.

Internet Search — Select which search engine to use for Internet searches. You can also change how the Search tab is used with searches.

Tabbed Browsing — Use these selections to have search results appear in tabs on the Mozilla Navigator window, as opposed to appearing in the full screen.

Download — Choose what you see during downloads from the Internet (a download manager, a progress dialog, or nothing).

The Advanced Preferences can be used to fine-tune your Web browsing experience, as well as make decisions on what content you open can and can't do on your computer. Here are some Advanced Preferences that might interest you:

Scripts & Plugins — The Web content you choose can try to open, move, resize, raise and lower windows. It can request to change your images, status bar text, or bits of information stored in what are called "cookies." These preferences let you restrict what the content you request can do.

Cache — By default, the most recent 4MB of Web pages you visit are stored in RAM and the most recent 50MB of pages you visit are stored on your hard disk. If a page is not out-of-date, caching makes it possible to return to a page quickly, without needing to reload from the original Web server. Cache preferences let you change how much information is cached, where hard disk cache is located. Other cache preferences let you clear all memory and disk cache immediately.

Proxies — If you have direct access to the Internet, you don't need to change any proxy settings. However, if you need to access the Internet via a proxy server, you can identify the location of that server (or servers) here. To access the Web via proxy servers, you must explicitly identify the proxy server to use for each type of content you request (HTTP, SSL, FTP, Gopher, and SOCKS).

HTTP Networking — If you are using a proxy server, that proxy server may

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

require that you make HTTP requests using either HTTP/1.0 or HTTP/1.1 standards (1.1 is the default).

Software Installation — When a request for Web content results in that content trying to install software on your computer, you can either allow that to happen (with prompting) or not allow any software updates, based on how this preference is set.

Mouse Wheel — If your mouse has a wheel, you can change how Mozilla behaves when you use that wheel. By default, the mouse wheel scrolls a line at a time.

DOM Inspector — Turn this on to check the structure of a Web page (for debugging).

Adding helper apps

Although the main type of content provided by Web pages is HTML, many other content types can be displayed, played, or presented by a Web browser. Most additional data encountered by Mozilla is handled in one of two ways: plug-ins or helper apps.

Plug-ins are self-contained programs that allow data to play within the Mozilla window. A helper app can be any program that is available on your Red Hat Linux system. It is up to you to identify the plug-in or helper app to launch when a certain type of data is encountered.

Mozilla determines what helper app or plug-in to launch based on the following criteria:

Suffixes — If the browser is reading a file that has a particular suffix attached to the filename (such as `exe` for an application or `gz` for a compressed zip file), it can use that suffix to determine the file's contents. When a file's suffix matches a suffix configured for a particular helper app or plug-in, the helper app or plug-in is used to play or display the data.

MIME type — Because data may come to the browser in a stream or have no suffix, Mozilla can use the MIME type attached to the data to determine which plug-in or helper app to use. (MIME stands for Multipurpose Internet Mail Extensions.)

You can add your own helper app to automatically handle a particular type of data in your browser. Here's how:

1. Choose Edit → Preferences. The Mozilla Preferences window appears.
2. Click the plus next to the Navigator category, and select Helper Applications.
3. Click New Type. A dialog box appears that enables you to add information about the helper app and the data that it can handle.
4. Type in a description of the data, the MIME type, and the file suffixes (if any) on files that contain that type of data.

   Note When you add the suffix, don't include the dot.

5. Choose an application to handle the data type. If the application needs a Terminal window to run, type **xterm -e**, followed by the command line you need to enter. This executes the command in an xterm window, reading in data as

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

needed.

6. Click OK when you are done.

The next time you open data in Mozilla of the type you just added, a pop-up window will ask you if you want to use the application you just entered to open the data.

Adding plug-ins

There are not many plug-ins available for use in the version of Mozilla that comes with Red Hat Linux. To see a list of plug-ins associated with your Mozilla browser, choose Help → About Plug-ins. You will probably see one plug-in: Plugger. If Plugger isn't installed, you can install it from CD #3. Plugger (`http://fredrik.hubbe.net/plugger.html`) is a multimedia plug-in that handles QuickTime, MPEG, MP2, AVI, SGI-movie, Tiff, DL, IFF-anim, MIDI, Soundtracker, AU, WAV, and Commodore 64 audio files by relying on external programs to play the data. Recently, Plugger has also added support for displaying a variety of Microsoft file types.

Here is a list of some free plug-ins available for Mozilla that work with Red Hat Linux.

Adobe Acrobat Plug-in (`www.adobe.com/support/downloads`) — Displays PDF (Portable Document Format) files from Adobe Systems.

Cult3D Plug-in (`www.cult3d.com`) — Displays high-quality, interactive real-time 3D images on the Web. This plug-in is from Cycore Computers. (Click Download to find the plug-in. There is a cult3d RPM available that requires tailoring to work in Mozilla.)

DjVu Plug-in (`djvu.sourceforge.net`) — Displays images in DjVu image compression technology. This plug-in is from AT&T.

Real Audio (`www.real.com`) — Plays Real Audio and Video content. (To use the Real Play plug-in in Mozilla, you need the RealPlayer and realplay-nplugin RPMs, available by searching `rpmfind.net`).

Macromedia Flash Player (`www.macromedia.com/`) — Displays multimedia vector graphics and animation. This plug-in is from Macromedia, Inc. The Flash Player is available for Red Hat Linux, but Macromedia's Shockwave plug-in is not yet available for Linux.

Xswallow (`www.csn.ul.ie/~caolan/docs/Xswallow.html`) — Allows X window programs to be used as an inline viewer.

More plug-ins may be added in the future for use with Red Hat Linux. Mozilla now offers a great new Web site for getting plug-ins for Linux that run in Mozilla: `http://plugindoc.mozdev.org/linux.html`. The site contains links to download plug-ins and instructions for installing them.

When you download a plug-in, follow the instructions that come with the plug-in for installing it. If the plug-in comes in an RPM file, install it as you would any other software package in Red Hat Linux (`rpm -iv` *`package`* command). Otherwise, probably just copy the plug-in file (a `.so` file) to the system plug-in directory (probably `/usr/lib/mozilla/plugins`) or your personal plug-ins directory (probably `$HOME/.mozilla/plugins`). When you restart Mozilla, the plug-ins will automatically be picked up from those locations.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The CrossOver Plugin is a commercial product that lets you use many Windows plug-ins in Mozilla. See the section "Adding a CrossOver Plugin" later in this chapter for

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using Mozilla Navigator controls

If you have used a Web browser before, the controls are probably as you might expect: location box, forward/back buttons, file and edit menus, and so on. There are a few controls with Mozilla, however, that you might not be used to seeing. Here are some examples:

Using the sidebar — Press the F9 function key to toggle the sidebar on and off. Click a What's Related, Search, Bookmarks, or History tab (or add your own tab as described earlier). Add your own bookmarks, return to pages from your history list, or search for Web content.

Sending Web Content — You can send an e-mail containing either the current Web page (File → Send Page) or the URL of the current Web page (File → Send Link) to selected recipients.

Searching the Internet — You can search the Internet for a keyword phrase in many different ways. Choose Tools → Search the Web to open a Netscape Web site that lets you search the Internet. Or type one or more keywords in the Location box, and click Search. Then, of course, you can use the sidebar for searches as described earlier.

Viewing Information about a Page — You can view information about the location of a Web page, the location of each of its components, the dates the page was modified, and other information by choosing View → Page Info. Figure 9-4 shows the Page Info window for a Web page. Click the Links tab to see links on that page to other content on the Web. Click the Security tab to see information about verification and encryption used on the page.

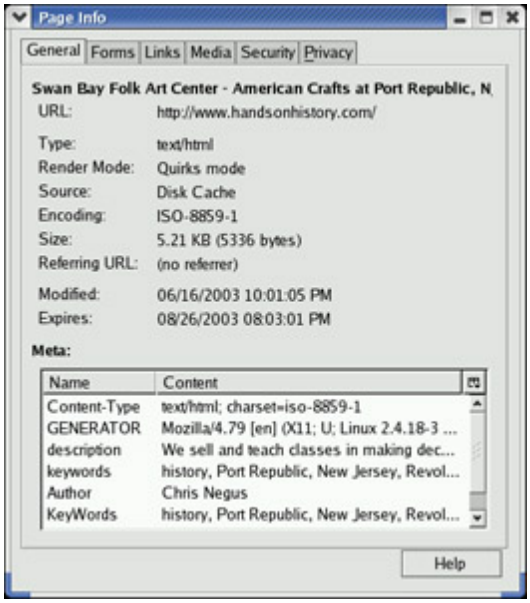

Figure 9-4: Display information about a Web page by selecting Page Info.

Improving Mozilla browsing

Every Web site you visit with Mozilla is not going to play well. Some sites don't follow standards, use unreadable fonts, choose colors that make it hard to see, or demand that you use a particular type of browser to view their content. To improve your browsing experience, there are several things you can add to Mozilla.

Adding a CrossOver Plugin

QuickTime 5 movies, Shockwave Director multimedia content, and various Microsoft movie, file, and data formats simply will not play natively in Mozilla Navigator. Using software built on WINE, CodeWeavers (www.codeweavers.com) created the CrossOver Plugin. Although it costs a few dollars ($24.95 U.S. for a one-user-at-a-time type license), the CrossOver Plugin lets you play some content that you simply could not otherwise use in Linux.

After you install the CrossOver Plugin, you see a nice Plugin Setup window that lets you selectively install plug-ins for QuickTime 6, Windows Media Player 6.4, Shockwave 8.5, Flash 6, and Microsoft

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Word, Excel, and PowerPoint viewers. You can also install other multimedia plug-ins, as well as a variety of fonts to use with those plug-ins.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Although I have had Cross Over Plugin running well in earlier versions of Red Hat Linux, at the time of this writing it was not yet working in the current version of Red Hat Linux. Check wit

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Adding a Preferences toolbar

Did you ever run into a Web page that required you to use a particular type or version of a browser or had fonts or colors that made a page unreadable? The Mozilla Preferences Toolbar lets you try to spoof Web sites into thinking you are running a different browser. It also lets you choose settings that might improve colors, fonts, and other attributes on difficult-to-read pages.
You can install the neat little toolbar from the Mozdev.org site (`prefbar.mozdev.org`). Click the Install link; then, after it is installed, restart Mozilla. Figure 9-5 shows an example of the Preferences Toolbar that has been installed in Mozilla.

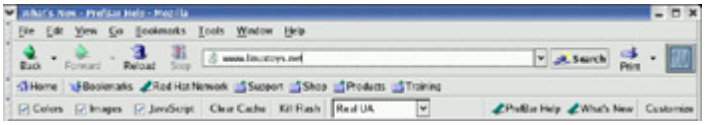

Figure 9-5: Change colors, fonts, and browser types on the fly with Mozdev.org Preferences Toolbar.

The default set of buttons let you do the following:

Colors — Change between default colors and those set on the Web page.
Images — Toggle between having images loaded or not loaded on pages you display.
JavaScript — Allow or disallow JavaScript content to play in Mozilla Navigator.
Clear Cache — Deletes all cached content from memory and disk.
Kill Flash — Kills all embedded Flash content on the current page.
Real UA — Choose to have your browser identified as itself (current version of Mozilla) or any of the following: Mozilla 1.0 (in Windows 98), Netscape Navigator 4.7 (in Macintosh), Netscape 6.2 (in Linux), Internet Explorer 5.0 (in Macintosh), or Internet Explorer 6.0 (in Windows XP)

Click the Customize button to add other buttons to the Preferences Toolbar. You can add buttons to clear your History or Location bar entries. You can even add a Popups button to prevent a page from opening a pop-up window from Mozilla.

Many of the preferences take effect immediately. Others may require you to restart Mozilla.

Adding Java support

If you want to display some Java content, but you only see a broken puzzle piece and a failure message that says you need a plug-in to view application/x-java-whatever content, you can install the software you need from Sun Microsystems (`www.sun.com`). Here's how:

1. Download the Java 2 Runtime Environment package from `java.sun.com/download`. (Look for a Java 2 Platform, Standard Edition that runs on Linux, then download the Java 2 Runtime Environment package. I used `j2re-1_4_2-linux-i586-rpm.bin`.)
2. Make the package executable and execute it to extract the RPM, as follows:
3. # **chmod 755 j2re*-rpm.bin**
4. # **./j2re*-rpm.bin**
5. Install the j2re3 package (as root user), just as you would any software RPM package.
6. # **rpm -iv j2re*rpm**
7. Change to the java plug-in directory and create a link from the java plug-in to the Mozilla plug-ins directory. (Instead of *, you could type the whole j2re directory name.)
8. $ **cd /usr/java/j2re*/plugin/i386/ns610/**
9. $ **ln -s libjavaplugin_oji.so /usr/lib/mozilla/plugins/**
10. Close any open Mozilla windows, then open a new Mozilla window.
11. Click Help → About Plug-ins. You should see a listing for the Java plug-in and lots of supported Java MIME types.
12. To see if the Java plug-in is working, you could try out some games from the Java Software Showcase (`http://java.sun.com/getjava/showcase.html`).

Doing cool things with Mozilla

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

There are some neat bells and whistles built into Mozilla that can make your browsing more pleasant. Here are a few of those features that you can try out:

Using tabbed browsing

If you go back and forth among several Web pages, you can use the tabbed browsing feature to hold multiple pages in your browser window at once. You can set up tabbed browsing to work from a Web page or from the Location box. Here's how:

1. Click Edit → Preferences. The Preferences window appears.
2. Click Tabbed Browsing under the Navigator category.
3. Click one or both of these boxes, depending on how you want to use tabbed browsing:
   Middle-click, Control+click or Control+Enter on links in a Web page — Selecting this box lets you open a link to another Web page in a new tab. For this to work, click the middle mouse button on a link, hold the Ctrl key while you click the left mouse button on a link, or (with the link highlighted) hold the Ctrl key and press Enter.
   Control+Enter in the Location bar — After you type a Web address (URL) into the Location box, hold the Ctrl key and press Enter to open the new page in a tab.
4. Click OK. You can begin using the tabbed browser feature from your Mozilla Navigator.

A tab for each tabbed page appears at the top of the Navigator pane. To close a tab, create a new tab, bookmark a group of tabs, or reload tabs, right-click one of the tabs at the top of the pane. A drop-down menu lets you choose the function you want.

Using the DOM Inspector

If you are debugging a Web page that you are creating, the Document Object Model (DOM) Inspector can be useful for checking out the structure of your page. To open the DOM inspector, from the Mozilla window click Tools → Web Development → DOM Inspector.

From the DOM Inspector window, type the URL to the Web page you want to check out. The nodes, representing the head, body, tables, fonts, and so on, appear in the left column. Values for each node appear in the right column. Click a node name and the selected area is highlighted on the page below, with the node value appearing to the right.

Resizing the Web page

There is a nice keyboard shortcut that lets you quickly resize the text on most Web pages in Mozilla. Hold the Ctrl key and press plus (+) or minus (-) keys. The text on the Web page will (in most cases) get larger or smaller, respectively. That page with the insanely small type font is suddenly readable.

Using text-based Web browsers

If you become a Linux administrator or power user, over time you will inevitably find yourself working on a computer from a remote login or where there is no desktop GUI available. At some point while you are in that state, you will probably want to check an HTML file or a Web page. To solve the problem, Red Hat Linux includes several text-based Web browsers.

With text-based Web browsers, any HTML file available from the Web, your local file system, or a computer where you're remotely logged in can be accessed from your shell. There's no need to fire up your GUI or read pages of HTML markup if you just want to take a peek at the contents of a Web page. Besides letting you call up Web pages, move around with those pages, and follow links to other pages, some of these text-based browsers even display graphics right in a Terminal window!

Which text-based browser you use is a matter of which you are more comfortable with. Browsers that are available include:

links — With `links`, you can open a file or a URL, and then traverse links from the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

pages you open. Use search forward (`/string`) and back (`?string`) features to find text strings in pages. Use up and down arrows to go forward and back among links. Then press Enter to go to the current link. Use right and left arrows to go forward and back among pages you have visited. Press Esc to see a menu bar of features to select from.

lynx — The `lynx` browser has a good set of help files that come with it (press the ? key). Step through pages using the Spacebar. Though `lynx` can display pages containing frames, it cannot display them in the intended positioning. Use arrow keys to display the selected link (right arrow), go back to the previous document (left arrow), select the previous link (up arrow), and select the next link (down arrow).

w3m — The `w3m` text-based Web browser can display HTML pages containing text, links, frames, and tables. It even tries to display images (though it is a bit shaky). There are both English and Japanese help files available (press H with `w3m` running). You can also use `w3m` to page through an HTML document (for example, `cat index.html | w3m`). Use Page Up and Page Down to page through a document. Press Enter on a link to go to that link. Search forward and back for text using / and ? keys, respectively.

The `w3m` command seems the most sophisticated of these browsers. It features a nice default font selection, seems to handle frames neatly, and its use of colors also makes it easy to use. The `links` browser lets you use the mouse to cut and paste text.

You can start any of these text-based Web browsers by giving it a filename, or if you have an active connection to the network, a Web address. For example, to read the w3m documentation (which is in HTML format) with a `w3m` browser, you can type the following from a Terminal window or other shell interface:

$ **w3m /usr/share/doc/w3m-0*/doc/MANUAL.html**

An HTML version of the W3M Manual is displayed. You can also have started links without giving it a filename (in which case, you will have to open one later from the menu). Or you can give w3m a URL to a Web page, such as the following:

$ **w3m www.handsonhistory.com**

After a page is open, you can begin viewing the page and moving around to links included in the page. Start by using the arrow keys to move around and select links. Use Page Up and Page Down keys to page through text.

Team LiB

PREVIOUS NEXT

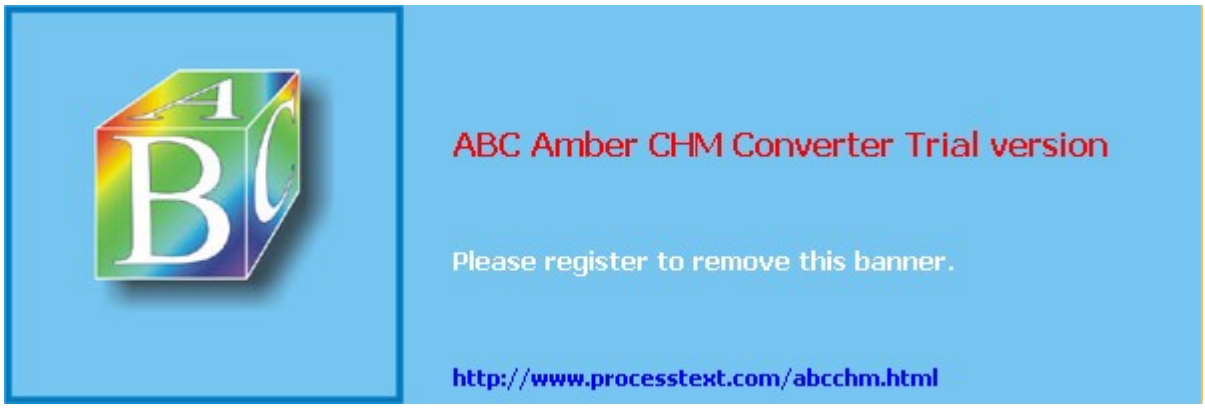

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Communicating with E-mail

To manage your e-mail, Red Hat Linux features the Ximian Evolution mail program. Mozilla also includes a graphical application for reading, managing, and sending e-mail.
If you don't mind text-based interfaces, or if you are a UNIX person who likes to `sort`, `grep`, `troff`, `col`, and `cat` your e-mail, there are still plenty of UNIX-like mail tools around. The `mail` command itself provides an easy interface for plain-text messages sent to other users on your UNIX system or on your LAN. There are also text-based mail applications, such as the `mutt` command, that let you handle mail attachments.

After covering some e-mail basics, this section will lead you through the steps that allow you to use e-mail with Evolution and Mozilla Mail. If you are interested in text-based, command-driven mail tools, some of which have been around UNIX systems for many years, descriptions of many of those commands are also in this section.

E-mail basics

E-mail is one of the oldest uses of computer networks — predating the Web by more than 20 years. In fact, e-mail was one of the first applications used to transport information on the Internet, when the Internet consisted of only a few computers.
Today, there are millions of users around the world who have e-mail addresses. Although there are several different styles of e-mail addressing, by far the most popular e-mail address format is the domain style address (used with the Internet and other TCP/IP networks). The e-mail address consists of a user name and domain name, separated by an `@` sign. For example:

webmaster@handsonhistory.com

As someone using e-mail, you need a program (such as Ximian Evolution) that enables you to get your e-mail, manage your e-mail messages, and send messages. Although mail messages were originally only plain text, and still are in most cases today, there are some newer features that let you enhance the kinds of content that you can send and receive. Here are two ways to enhance your mail messages:

Attachments — You can attach files to your mail messages. Attachments can contain data that you couldn't ordinarily keep in a mail message, such as a binary program, a word processing file, or an image. The recipient of the mail attachment can either save the file to a local hard disk or open it in a program designed to read the attachment.

HTML — The same stuff used to create Web pages can be included in mail messages you create with certain mail clients (including Evolution and Mozilla Mail). This lets you change fonts and colors, add backgrounds, insert images, or add HTML features.

Caution To people who use text-based mail clients, HTML content can't be interpreted (it shows up as a bunch of markers that overwhelm the text). In general, don't use HTML in messages that are being distributed to a large group of people (such as in a newsgroup). Also, e-mail was never intended to transport large attachments. For larger files, try copying to an FTP site instead of doing e-mail attachments.

Depending on the mail program you are using (which is also referred to as a Mail User Agent or MUA), e-mail management features let you direct incoming e-mail into different folders and sort messages by date, sender, or other attributes. E-mail sending features let you reply to messages, forward messages, and draw names from an address book or directory server.

Cross-Reference If you don't have an e-mail account, you can set up your own e-mail server using Red Hat Linux. For information on setting up a mail server, see Chapter 19.

Using Evolution e-mail

Evolution is the preferred application for sending and managing e-mail in Red Hat Linux. Red Hat gave it a prime spot on the desktop, just to the right of the Red Hat main menu and Web browser icons. After you launch Evolution for the first time and run the Startup Assistant, the Evolution window appears,

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

showing the different types of operations you can perform.
Figure 9-6 shows an example of the Evolution window. Evolution is a groupware application, combining several types of applications that help groups of people communicate and work together. The features of Evolution include:

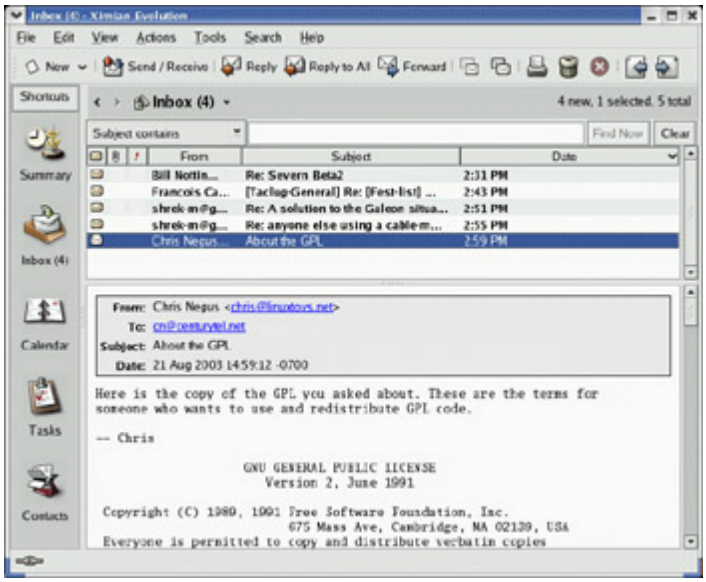

Figure 9-6: Evolution can be used to manage your mail, appointments, and tasks.

E-mail — Includes a complete set of features for getting, reading, managing, composing, and sending e-mail on one or more e-mail accounts.

Calendar — Create and manage appointments on your personal calendar. You can e-mail appointment information to others and do keyword searches of your calendar.

Tasks — Organize ongoing tasks into folders.

Contacts — Create contact information for friends and associates, such as names, addresses, and telephone numbers. A Categories feature helps you remember who gets birthday and anniversary gifts.

In the next section, I focus on the Summary and E-mail features of Evolution.

Setting Evolution preferences

To really make Evolution your own, you can set preferences that are particular to you. Certain settings let you choose the information that appears on the Summary page when you start Evolution. Other settings let you choose how your e-mail is gathered and sorted.

Changing Summary settings

You can change Summary settings by performing the following steps:

1. From the Evolution main window, select Tools →Settings.
2. Select Summary Preferences (left column). You can change the following Summary settings:

Mail — Select which folders should appear on the Summary window. By default, only the Mail Summary link appears, but no specific folders are shown.

News Feeds — Select from a variety of news feeds, to have headlines and links to news stories added to your Summary page. Click the News Feeds tab, choose the news headlines that interest you, and click Add to have them appear on the list. By default, news headlines will refresh every 10 minutes (600 seconds), and up to 10 headlines will appear for each feed (both values are configurable).

Weather — Choose to have weather for a particular city appear on your Summary page. Click the Weather tab, select the country, state, and city for which you want to see weather, and then click Add to make the change.

Schedule — Choose how many days of scheduled appointments

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

to appear on your Summary page. Click the Schedule tab and select one day, five days, one week, or one month of appointments. Also, select to show all or just today's tasks.

3. Click Apply to apply the changes.

Changing Mail Accounts settings

You can change Mail Accounts settings by performing the following steps:

1. From the Evolution main window, select Tools →Settings.
2. Click Mail Accounts in the left column.
3. Select the mail account to change and click Edit. The Evolution Account Editor appears.
4. Here are a few items you may want to change for your e-mail account:
   Signature — Have a signature appear on every e-mail message you send. Either click the Default signature box and select Autogenerated (to use name and e-mail address as a signature) or click Add New Signature to create a signature.
   Receiving — By default, Evolution checks your mail server for your mail every ten minutes. Once downloaded, each message is erased from the server. To change automatic e-mail checking and options to leave messages on the server, select the Receiving Options tab and make the changes you want.
   Automatic copy — You can have every message you send copied to one or more other users. This is a nice feature if you write important e-mail that you want to archive to a different e-mail account. Select the Defaults tab, then click the check box next to Always Cc or Always Bcc. Next, type the correct e-mail address.
   Security — To help validate that you are who you say you are and keep your e-mail private, Evolution lets you use PGP (Pretty Good Privacy) encryption keys. Click the Security tab, then enter your PGP/GPG Key ID. Choose settings for signing and encryption as appropriate.
5. Click Apply to apply the changes.

Receiving, composing, and sending e-mail

Evolution offers a full set of features for sending, receiving and managing your e-mail. I personally prefer the Folder to the Shortcut view for working with Evolution (click View → Folder Bar to turn it on; click View → Shortcut Bar to turn it off).

Here are some tips for sending, reading, and receiving mail:

Read e-mail — Click Inbox in the Folder column. Your messages appear to the right.

Delete e-mail —After you have read a message, select it and press the Delete key. Click View → Hide Deleted Messages to toggle whether or not you see deleted messages. Click Actions → Expunge to permanently remove all messages marked for deletion in the current folder.

Send and receive — Click the Send/Receive button to send any e-mail queued to be sent and receive any e-mail waiting for you at your mail server.

Compose e-mail — Click New → Mail Message. A Compose a Message window appears. Type the e-mail address, a message for the subject line, and the body of the message. Click Send when you are finished. Buttons on the Compose window let you add attachments, cut and paste text, choose a format (HTML or plain text), and sign the message (if you have set up appropriate keys).

Create folders — If you like to keep old messages, you may want to save them outside your Inbox (so it won't get too junked up). Right-click on the Inbox, then select New Folder. Type a folder name and click OK (to store it as a subfolder to your Inbox).

Sort messages — With new folders created, you can easily sort messages from your Inbox to another folder. The easiest way is to simply drag-and-drop each message (or a set of selected messages) from the message pane to the new

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

folder.

Search messages — Type a keyword in the search box over your e-mail message pane and select whether to search your message subject lines, sender, recipient, or message body. Click Find Now to search for the keyword. After viewing the messages, click Clear to have the other messages reappear.

Filter messages — You can take action on an e-mail message before it even lands in your Inbox. Click Tools →Filters. A Filters window appears that lets you add filters to deal with incoming or outgoing messages. Click Add to create criteria and set actions.

For example, you could have all messages from a particular sender, subject, date, status, or size sorted to a selected folder. Or you could have messages matching your criteria deleted, assigned a color, or respond by playing a sound clip.

Cross-Reference Refer to Chapter 19 for information on using Spamassassin, along with Evolution filters, to sort out SPAM from your real e-mail messages.

Besides the features mentioned in the previous list, Evolution supports many common features, such as printing, saving, and viewing e-mail messages in various ways. The help system that comes with Evolution (click the Help button) includes a good manual, FAQ, and service for reporting bugs.

Mozilla Mail client

The Mozilla Mail client program provides a simple interface for using e-mail in Red Hat Linux. You can access the Mozilla Mail window from your Mozilla browser window by choosing Tasks → Mail & Newsgroups. Figure 9-7 shows an example of the Mozilla Mail window that is ready to use mail and news.

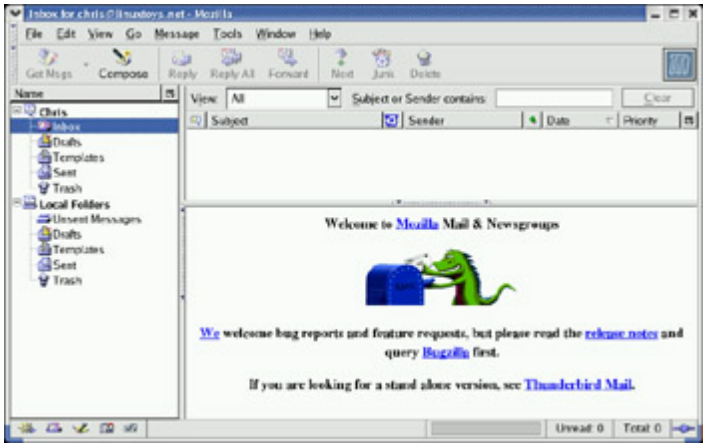

Figure 9-7: Manage your e-mail from the Mozilla Mail window.

As with Evolution e-mail, you need to provide some information about your mail account before you can connect to your mail server and use Mozilla Mail.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip A new Junk Mail feature was recently added to Mozilla Mail. With Junk Mail, Mozilla Mail automatically tags any message it believes to be junk mail with a blue recycl

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Connecting to the mail server

You must identify information about yourself and your mail server before you can download or send mail messages. The first time you open Mozilla Mail, an Account Wizard appears. Select to set up an e-mail account; then enter the following information as you are prompted for it:

Your Name — The name to appear on messages you send (for example, John Smith).
Your E-mail Address — The *name@domain* address that is assigned to you.
Incoming Mail Server — The name or IP address of the mail server from which you get your messages. You also need to know what type of server it is (POP or IMAP). POP servers require a POP login and password. POP stands for Post Office Protocol. IMAP stands for Internet Message Protocol.
Outgoing Mail Server — If there is a different outgoing mail server, you need the name or IP address of that server. Likewise, you will probably need login and password information to send mail on that server. You can choose to use Secure Socket Layer (SSL) to protect your outgoing messages (if it is supported by the server).

Once the wizard has created the account, you are prompted for the password for your account on the mail server. Using that password, Mozilla Mail will download all your messages from the mail server. It will try to download messages again every 10 minutes. Or, you can click the Get Msgs button to download messages immediately.

If you want to change how often mail is downloaded, or other features of your account, click Edit → Mail & Newsgroup Account Settings. Under the e-mail account you added are categories to change the set-up and behavior of the account. (Click Server Settings to change how often, if at all, new messages are automatically downloaded from the mail server.)

Managing incoming mail

There are various ways to store and manage the e-mail messages in Mozilla Mail. The following is a quick rundown of how to use Mozilla Mail for managing incoming mail:

Your e-mail messages are typically stored on a mail server that is contacted over the network. To download your mail messages immediately from the mail server, click File → Get New Messages (or click Get Msgs on the toolbar).
Mail messages are stored in folders under the Mail Folders heading in the left column. There should be a separate heading for each mail account you have. For each mail account, incoming messages are stored (by default) in your Inbox folder. You can create additional folders to better keep track of your mail (right-click on Inbox, then select New Folder to add a folder). Other folders contain drafts of messages set aside for a time (Drafts), templates for creating messages (Templates), messages you have sent (Sent), and messages that you have discarded (Trash).
Messages are sorted by date for the folder you select, in the upper-right corner of the display. Click on the heading over the messages to sort by subject, sender, or priority.
When you select a message, it appears in the lower-right corner of the display.

Composing and sending mail

To compose e-mail messages, you can either start from scratch or respond to an existing e-mail message. The following are some quick descriptions of how to create outgoing mail:

To create a new message, choose File → New → Message (or Compose on the toolbar).
To reply to a mail message, click on the message on the right side of your screen, and then choose Message → Reply (to reply only to the author of the message) or Message → Reply to All (to reply to everyone listed as copied on the message).
To forward a mail message, choose Message → Forward. You can also forward a message and have it appear in the text (Message → Forward As → Inline) or as an attachment (Message → Forward As → Attachment).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

In each case of outgoing mail, a mail Compose window appears, in which you compose your e-mail message. As you compose your message in the Compose window, you can:

- Add e-mail addresses from your personal address book (or from one of several different directory servers) by choosing Options → Select Addresses. A list of your stored addresses appears for you to choose from. Click Collected Addresses to see a list of addresses that have been collected from e-mail messages you have received.
- Add attachments to the message (such as a word processing file, image, or executable program) by choosing File → Attach File. After that, you can select a file from your file system to attach. (You can also choose File → Attach Web Page to choose the URL of a Web page that you want to attach.)
- Add certificates or view security information about your mail message by clicking View →Message Security Info. In the window that appears, you can view general security information about the message.

When you are finished composing the message, click Send to send the message. If you prefer, queue the message to be sent later by choosing File → Send Later. (Send Later is useful if you have a dial-up connection to the network and you are not currently online.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you want to quit and finish the e-mail message later, you can choose File → Save As → Draft. Then click the X in the upper-right corner to close the window. When you are read

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Text-based mail programs

There are many text-based mail programs for reading, sending, and working with your mail. Many of these programs have been around for a long time, so they are full of features and have been well debugged. As a group, however, they are not very intuitive. The following sections describe some text-based commands.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Most of these programs use the value of your $MAIL environment variable as your local mailbox. Usually, that location is /var/spool/mail/*user*, where *user* is your user name. To set your $MAIL so that it points to your Mozilla mailbox (so you can use either Mozilla Mail or a text-based mail program), add the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Mail readers and managers

Mail readers described below are text-based and use the entire screen. Although some features are different, menu bars show available options right on the screen.

Mutt mail reader

Note To use the mutt mail reader you must have the mutt software package installed from CD #1.

The `mutt` command is a text-based, full-screen mail user agent for reading and sending e-mail. The interface is quick and efficient. Type **mutt** to start the mail program. Move arrow keys up and down to select from your listed messages. Press Enter to see a mail message and type **i** to return to the Main menu.

The menu bar indicates how to mark messages for deletion or undelete them, save messages to a directory, or reply to a message. Type **m** to compose a new message and it opens your default editor (for me, `vi`) to create the message. Type **y** to send the message. If you want to read mail without having your fingers leave your keyboard, `mutt` is a nice choice. (It even handles attachments!)

Pine mail reader

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note To use the pine mail reader, you must have the pine software package installed. The pine package is no longer distributed with Red Hat Linux. You can try rpmfind.net to get a pin

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The `pine` mail reader is another full-screen mail reader, but it offers many more features than does `mutt`. With `pine`, you can manage multiple mail folders. You can also manage newsgroup messages, as well as mail messages. As text-based applications go, `pine` is quite easy to use. It was developed by a group at the University of Washington for use by students on campus, but has become widely used in UNIX and Linux environments.
Start this mail program by typing `pine`. The following menu is displayed, from which you can select items by typing the associated letter or using up and down arrows and pressing Enter:

? HELP - Get help using Pine

C COMPOSE MESSAGE - Compose and send a message

I MESSAGE INDEX - View messages in current folder

L FOLDER LIST - Select a folder to view

A ADDRESS BOOK - Update address book

S SETUP - Configure Pine Options

Q QUIT - Leave the Pine program

To read your e-mail, select either `I` or `L`. Commands are listed along the bottom of the screen and change to suit the content you are viewing. Left (() and right (→) arrow keys let you step backward and forward among the `pine` screens.

Mail reader

The `mail` command was the first mail reader for UNIX. It is text-based, but not screen-oriented. Type **mail** and you will see the messages in your mailbox. Because `mail` is not screen-oriented, you just get a prompt after message headings are displayed — you are expected to know what to do next. (You can use the Enter key to step through messages.) Type **?** to see which commands are available. While in `mail`, type **h** to see mail headings again. Simply type a message number to see the message. Type **d#** (replacing # with a message number) to delete a message. To create a new message, type **m**. To respond to a message, type **r#** (replacing # with the message number).

Team LiB

PREVIOUS NEXT

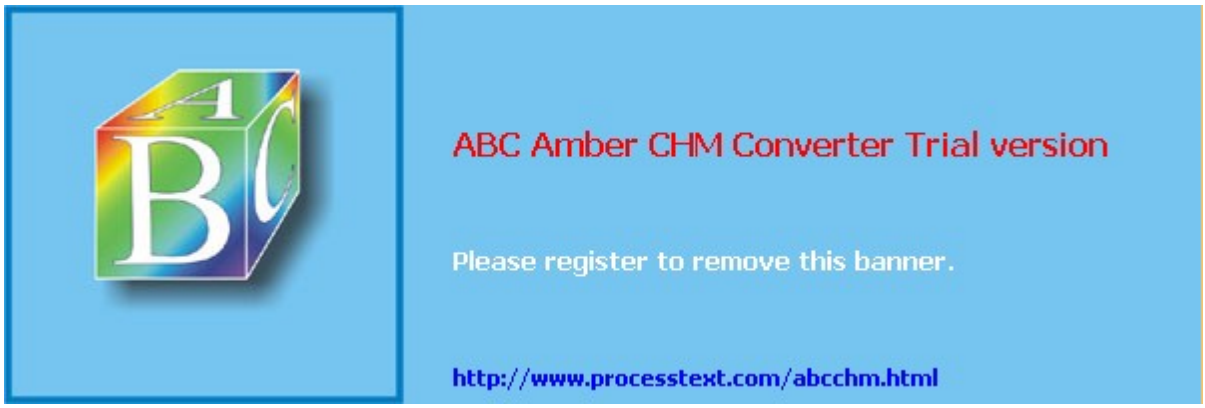

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Participating in Newsgroups

Usenet news is another feature that has been around almost as long as the Internet. Using a newsreader, and even many regular mail readers, you can select from literally thousands of topics and participate in discussions on those topics. To participate, you simply read the messages people have posted to the group, respond to those that you have something to say about, and send your own messages to start a discussion yourself.

To get started, you basically need a newsreader and access to a news server computer. As with e-mail, Mozilla uses Mozilla Mail to let you participate in newsgroups. Another popular newsreader that comes with Red Hat Linux is called pan.

Tip If you have never used a newsgroup before, check out the news.announce.newusers newsgroup. It's there to answer questions from new users.

Mozilla for newsgroups

The same Mozilla Mail window you use for e-mail can be used to participate in newsgroups. Before you can begin using newsgroups, you need to identify the news server you are using, identify yourself, and choose the newsgroups you want to access.

Connecting to the news server

To begin using newsgroups, open the Preferences window from Mozilla and add the news server's name and port number as follows:

1. From the Mozilla Navigator window, click Window → Mail & Newsgroups.
2. From the Mail window, choose Edit → Mail & Newsgroups Account Settings to bring up the window to create the news server account.
3. In the Mail & Newsgroups Account Settings window, click Add Account. The Account Wizard window appears.
4. Click Newsgroup Account and click Next. You are asked to identify yourself.
5. Type your name (as you would have it appear to people in newsgroups) and your e-mail address (where you would like newsgroup participants to send you mail), and click Next.
6. Type the address of your news server (such as `news.handsonhistory.net`) and, if necessary, the port number for news on that server; then click Next. (The default news port is 119. You can obtain the name of the news server from your ISP or your system administrator at work. By convention, the name often has "news" or "nw" at the beginning.)
7. Type a name to represent the news account (it can be the server address), and click Next.
8. If the information you have entered looks correct, click Finish. The news server you have just added should now appear in your Mail & Newsgroups Account Settings window.
9. With the basic newsgroup settings in place, you can now set a variety of options to define how your newsgroup interface behaves. For example, you can change the following:

   Account Settings — Besides the Account Name, Your Name, and E-mail Address (which should already be there), you can identify a different Reply-to e-mail address, identify your organizaiton, and attach a signature. You can also have messages you send to newsgroups be in HTML format (usually you should not use HTML in newsgroups).

   Copies & Folders — You can choose to have messages you send to newsgroups copied to your Sent folder (default). You can also choose to have blind copies of the messages sent to yourself or another e-mail address.

   Addressing — If you have preferences set in an LDAP server, you can choose to have those preferences used with this news account.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Offline & Disk Space — You can limit the amount and size of news messages that are downloaded to your computer. You can also select how many messages are kept when you clean up your news messages.
Server Settings — For the news server, you can change the port number you connect to (usually 119 for news servers), whether or not to use a secure connection, how often you should check the server for messages, and whether or not you want to be prompted before downloading more than a set number of messages (500 by default).

10. Click OK.

If your news server was properly identified and your connection to that server is working, you can click on the news server's name in the Mozilla Mail window to begin participation in newsgroups. Connection to the news server occurs when you select news in that window. You must have a Mozilla mail account configured to post messages to a newsgroup.

Using newsgroups

To participate in newsgroups, you need to choose the ones that interest you from a list of available newsgroups. Before you start using newsgroups, create a valid mail account through Mozilla Mail (so you can fully participate once you begin). You can then subscribe to your favorite newsgroups so that they are available every time you use Mozilla.

Subscribing to newsgroups

The following procedure describes how to subscribe to a newsgroup in Mozilla.

1. Open Mozilla Mail & Newsgroups (Window → Mail & Newsgroups from Mozilla).
2. Right-click the news server name in the left column; then click Subscribe. A Subscribe window appears, as shown in Figure 9-8.

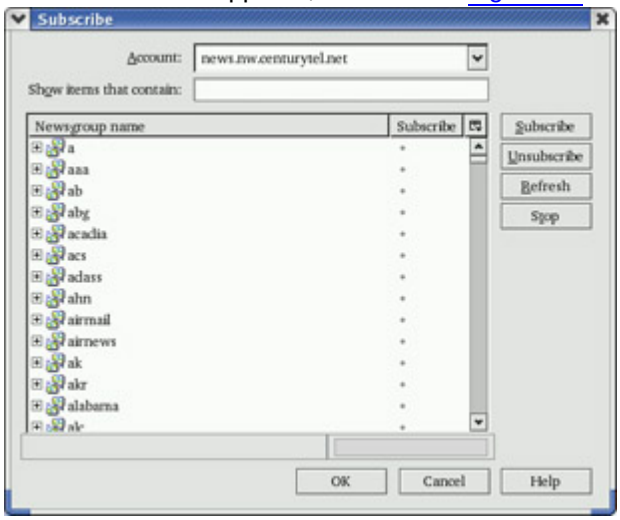

Figure 9-8: Choose from thousands of newsgroups in the newsgroup subscription window in Mozilla.

What appears (after a few minutes) are top-level newsgroup names. Newsgroups are organized in a hierarchy, similar to Internet domain names — levels are separated by dots (although newsgroup names start at the left and go down to the right).

3. Click the arrow sign next to a category to see which groups in that category interest you.
4. When you see a newsgroup that interests you, click it and select Subscribe.
5. Click OK. The subscribed newsgroup appears in your Mail & Newsgroups window under the name of the news server.
6. Click the newsgroup in the Mail & Newsgroups window. Recent messages from that newsgroup are downloaded to your Messenger window.

Reading newsgroup messages

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Reading newsgroup messages is similar to reading mail messages. You simply click on a message that interests you or move to it using the up and down arrow keys. The message will appear in the message pane in the lower-right corner of the Mail & Newsgroups window.
The concept of threads should be mentioned here: When someone posts a message to a newsgroup, usually one or more people respond to it. Because it is useful to read messages and responses together, these messages are grouped together in *threads*. A plus (+) next to a message indicates there are responses to that message. Click the plus to see available threads. Threads make it easy to see messages on a topic that interests you or skip a topic that doesn't.

The Pan newsreader

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note To use pan, you need to install the pan RPM package from CD #3 that comes with this book. For further information about pan, refer to: http://pan.rebelbase.com.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The Pan newsreader is a graphical application for reading, managing, and interacting with newsgroups. It is particularly adept at displaying attached pictures and downloading binaries. The interface is very intuitive and easy to use.

Pan is designed for the GNOME desktop. If you are running another desktop environment, however, you can still use Pan as long as you have the gnome-libs software package installed. After that, Pan works on KDE, Window Maker, or other desktop environments.

To open the Pan newsreader from the main Red Hat menu, choose Internet → More Internet Applications → Pan. The first time you start Pan, the Pan wizard runs to let you set up the newsreader. Have your e-mail address and your news server's name ready. When the wizard is done, you can download the list of newsgroups available from your news server.

Team LiB

PREVIOUS NEXT

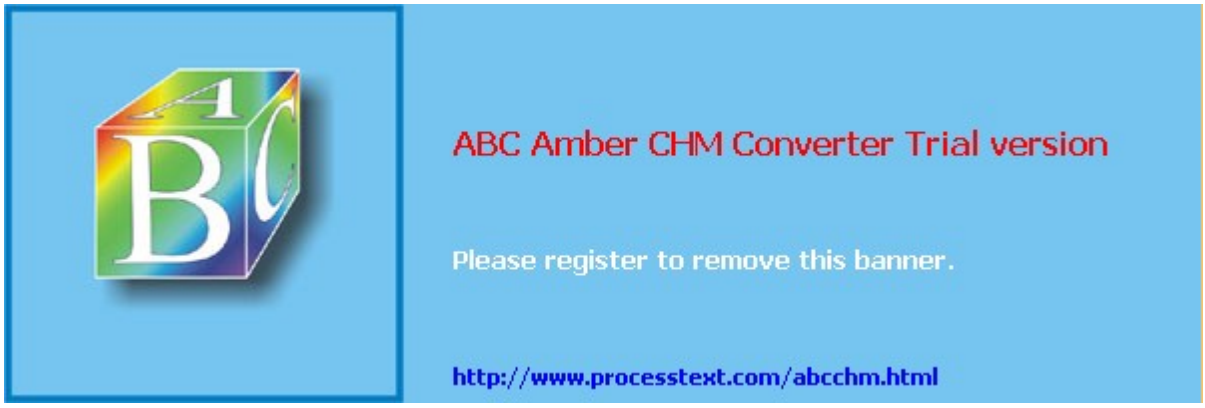

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Participating in AOL Instant Messaging with Gaim

Gaim is an AOL Instant Messaging client program that runs in Linux. Because it is based on AOL's Open IM architecture (http://www.aim.com), it has the same messaging features you get from AOL's own messaging application.

To start Gaim from a Red Hat desktop menu, choose Internet → Instant Messenger. Figure 9-9 contains an example of the Gaim window:

Figure 9-9: Access your AOL Instant Messaging using Gaim.

From the main Gaim window, you have the following options:

Accounts — Add screen names, passwords, and options associated with your AOL Instant Messaging accounts.

Signon — Sign on to AOL Instant Messaging by using your selected screen name.

Preferences — Set preferences associated with AOL Instant Messaging and Chat windows. You can also load plug-ins that can perform special functions, such as spell checking and auto-reconnection.

Besides AOL IM protocols, Gaim supports other messaging protocols, including instant messaging protocols from Yahoo!, the Microsoft Network, ICQ, and IRC.

ABC Amber CHM Converter Trial version

Please register to remove this banner.

http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using Remote Login, Copy, and Execution

This section describes some features for allowing users to use resources across a network. They are the `telnet`, `ssh`, `ftp`, `wget`, `rsync`, `rlogin`, `rcp`, and `rsh` commands.

Cross-Reference

Only the ssh service is turned on by default in Red Hat Linux. That is because the other remote login, execution, and copy commands described here do not provide encrypted communications by default, and so they can represent significant security risks. For information on how to turn on the services described in this chapter on a Red Hat Linux server, refer to Chapter 14 and Appendix C.

Two of the commands described in this section that are generally available are remote login and file transfer programs: `telnet` and `ftp`, respectively. Other commands for accessing FTP servers are `ncftp` and `gFTP`.

Two newer commands for copying files over the network are `wget` and `rsync`. Both of these tools can be used for efficiently downloading files that you can identify on the network.

Other commands described in this section are more specific to Red Hat Linux and other UNIX systems, but they provide simplified ways of copying files, logging in, and executing commands among trusted systems: `rcp`, `rlogin`, and `rsh`, respectively. Also, a more secure tool for remote login, called `ssh`, is included with Red Hat Linux. The `ssh` command can be used for remote execution as well.

Using telnet for remote login

Telnet is a service provided by many different types of computer systems to enable remote users to log in to their machines over TCP/IP networks. The `telnet` command is the client program that you use to do the remote login. The most common way to use telnet is with a host name. The following is a typical telnet session:

$ **telnet maple**

Trying 10.0.0.11 ...

Connected to maple (10.0.0.11).

Escape character is '^]'.

Red Hat Linux

Kernel 2.4.20-8 on an i686

login: **mike**

Password: **********

Last login: Mon Mar 31 13:15:57 from pine

[mike@maple mike]$

This example shows what happens when the `telnet` command is used to log in from a computer named `pine` to a computer named `maple` by typing `telnet maple`. My computer tries to connect to the telnet port on `maple` (IP address 10.0.0.11). Because `maple` is also a Red Hat Linux computer, once the connection is established, I see the standard `login:` prompt. I type the user name (`mike`) and the password when prompted. When the login and password are accepted, I see the shell prompt for the user named `mike`.

The telnet service is disabled by default on Red Hat Linux systems (as are most network services). So, to be able to log in to your computer using telnet, refer to Appendix C for information on how to turn on the telnet service.

Here are a few useful options you can use with telnet:

- **-a** — Automatic login. With this option, your computer attempts to log in to the remote computer using your local user name. So, if you are logged in to your computer as `mike`, when you use `telnet` to log in to a remote computer, the remote computer assumes that you want to log in as `mike`. It simply prompts you for `mike`'s password.
- **-l *user*** — User name. This option is similar to the `-a` option, except that instead of using your current user name, you can ask to log in using any user name you choose.
- **-r** — Rlogin-style interface. This option lets you use tilde (~) options. For example, to

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

disconnect while in rlogin mode, type **~.** (tilde+dot), or to suspend the telnet session, type **~^z** (tilde+Ctrl+z). Only use `~.` if your remote shell is hung (exit is a better way to quit normally). If you use `~^z` to temporarily suspend your telnet session, you are returned to your local system shell. To get back to the suspended session, type **fg** to put telnet back in the foreground. See the description of `rlogin` later in this chapter for more features you can use in rlogin mode.

Another way to use `telnet` is in command mode. Instead of using a host name, simply type the word **telnet**. You will see a `telnet` prompt as follows:

$ **telnet**

telnet>

At this point, there are several commands available to you. You are not yet connected to a remote host. To open a login session to a remote computer from the telnet prompt (for example, to a computer named `maple`), type:

telnet> **open maple**

After you do connect to a remote computer, you can return to the telnet session at any time by typing `Ctrl+]`. Here are other options you can use during your telnet session:

? — Print help information.
! — Escape to the shell.
`close` — If you have an open connection, type **`close`** to close it.
`display` — Shows the operating parameters that are in effect.
`logout` — Logs you off any remote connection in this session and closes it.
`mode` — Tries to enter line mode or character mode. (Type **`mode ?`** to see other options that go with the mode option.)
`quit` — Close telnet and exit.
`z` — Suspend the current telnet session.

If you suspend a telnet session or escape to the shell, return to telnet by typing `fg`.

Copying files with FTP

Like telnet, FTP is a protocol that is available on many different operating systems. Archives of files on the Internet are stored on what are called FTP servers. To connect to those servers from Red Hat Linux, you can either type the URL of that server into a Web browser or you can use the `ftp` command or graphical FTP windows such as gFTP. Of the other FTP client programs available with Red Hat Linux, my favorite is the `ncftp` command.

Using the ftp command

The `ftp` command is available on Red Hat Linux, as well as every other Linux and UNIX system, for copying files to and from FTP servers. Like telnet, FTP has a command mode or you (more typically) can use it to connect directly to a remote computer. For example:

$ **ftp maple**

Connected to maple.

220 (vsFTPd 1.1.3)

Name (maple:mike): **jake**

331 Please specify the password.

Password: **********

230 Login successful. Have fun.

Remote system type is UNIX.

Using binary mode to transfer files.

ftp>

In this example, `ftp` connects to a computer called `maple` (`ftp maple`). When I was prompted for a name, it assumed that I was going to use my current login name on `maple` (`maple:mike`). I could have pressed Enter to use the name `mike`, but instead I logged in as `jake` and typed the password when prompted. The password was accepted and, after some information was printed, I was given an `ftp>` prompt.

Because FTP is used for public servers, you can often log in using the word `anonymous` as your user

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

name. By entering a valid e-mail address as your password, you can enter the anonymous FTP site and download files that the server makes available to the public.

Unlike telnet, instead of being in a regular UNIX shell after I logged in with FTP, I was placed in FTP command mode. Command mode with FTP includes a whole lot of commands for moving around the remote file system and for copying files (which is its main job).

FTP directory commands

To get your bearings and move around the remote file system, you could use some of the following commands from the `ftp>` prompt. The commands are used to work with both the remote and local directories associated with the FTP connection.

pwd — Shows the name of the current directory on the remote system.
ls — Lists the contents of the current remote directory using the UNIX `ls` command. You can use any valid `ls` options with this command, provided that they are supported by the particular FTP server you are connected to.
dir — Same as `ls`.
cd — Use the `cd` command to move to the directory named.
cdup — Moves up one directory in the file system.
lcd — Lists the name of the current local directory.

If you want to make changes to any of the remote files or directories, use the following commands:

mkdir — Creates a directory on the remote system.
rename — Renames a file or directory on the remote system.
rmdir — Removes a remote directory.
delete — Removes a remote file.
mdelete — Removes multiple remote files.

Depending on how the FTP server is configured, you may or may not be able to do some of the file and directory commands shown above. In general, if you log in as the anonymous user, you will not be able to modify any files or directories. You will only be able to download files. If you have a real login account, you will typically have the same read and write permission you have when you enter the computer using a standard login prompt.

FTP file copying commands

Before you copy files between the remote and local systems, consider the type of transfer you want to do. The two types of transfer modes are:

binary — For transferring binary files (such as data files and executable commands). This is also referred to as an image transfer.
ascii — For transferring plain-text files.

The Linux `ftp` command seems to set the default to binary when you start FTP. Binary seems to work well for either binary or text files. However, binary transfers may not work transferring ascii files from non-UNIX systems. If you transfer an executable file in ascii mode, the file may not work when you try to run it on your local system.

Most file copying is done with the `get` and `put` commands. Likewise, you can use the `mget` and `mput` commands to transfer multiple files at once. Some FTP servers will even allow you to use matching characters (for example, `get abc*` to get all files beginning with the letters abc). Here are descriptions of those commands:

get file — Copies a file from the current directory on the remote file system and copies it to the current directory on the local file system. You can use a full path along with the filename. Here are some examples:

ftp> **get route**
ftp> **get /tmp/sting**

The first example takes the file `route` from the current remote directory and copies it to the current local directory. The second example copies the file `/tmp/sting` from the remote system to the file `tmp/sting` relative to the current directory on the local system. So if your current directory were `/home/jake`, `ftp` would try to copy the file to `/home/jake/sting`.

put *file* — Copies a file from the current local directory to the current remote directory. The usage of this command is essentially the same as the `get` command, except that files are copied from the local to the remote system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Anonymous FTP sites (described later) usually let you copy files *from* them, but not *to* them. If they do allow you to put files on their server, it will usually be in a restricted area.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

mget *file* . . . — This command lets you get multiple files at once. You can specify multiple files either individually or by using metacharacters (such as the asterisk). FTP prompts you for each file to make sure you want to copy it.
mput *file* . . . — This command lets you put multiple files on the remote computer. Like `mget`, `mput` prompts you before transferring each file.

FTP exiting commands

While a connection is open to a remote computer, you can use several commands to either temporarily or permanently exit from that connection. Here are some useful commands:

! — This temporarily exits you to the local shell. After you have done what you need to do, type **exit** to return to your FTP session. You can also use this command to run other local commands. For example, you can type **!pwd** to see what the current directory is on the local system, **!uname -a** to remind yourself of your local system name, or **!ls -l** to see the contents of your current directory.
close — Closes the current connection.
bye — Closes the connection and exits the `ftp` command.

Using the ncftp command

By virtue of being an FTP client program, the `ncftp` command supports all the standard commands you would expect to find in an FTP client (`get`, `put`, `ls`, `cd`, and so on). However, `ncftp` has added features that make it more efficient and more friendly than most FTP clients.
With `ncftp`, you can connect to an FTP server in the same way you did with `ftp`. One convenient difference is that if you enter no user name, `ncftp` assumes you want to use the anonymous user name and just logs you in. Here is an example:

$ **ncftp ns1**

NcFTP 3.1.5 (Oct 13, 2002) by Mike Gleason (ncftp@ncftp.com).

Connecting to 63.240.14.64...

(vsFTPd 1.1.3).

Logging in...

Login successful. Have fun.

Sorry, I don't do help.

Logged in to ns1.

ncftp / >
To log in as a user name other than anonymous, add a `-u` *`user`* option, where *user* is replaced by the name you want to log in as. Enter the password as prompted to continue.

Using ncftp

Once you are logged in, there are a few nice features you can use that aren't available with other FTP clients. Here are some examples:

bookmark — If you are visiting a site that you want to return to, type the bookmark command and type a name to identify that site. The next time you start an `ncftp` session, type the bookmark name as an option. Not only are you logged into the FTP site you bookmarked; you are taken to the directory where you set the bookmark.
lls, lcd, lmkdir, lpwd — There is a set of commands that let you move around the local file system. Just place a letter "`l`" in front of standard shell commands like `ls`, `cd`, `mkdir`, `pwd`, `rm` and `rmdir` and you can move around and work with your local file system while you are in `ncftp`.
rhelp, site — Use `rhelp` to see commands that are recognized by the remote FTP server. To see commands that are specific to the FTP server, type **help**.
Auto-resume — If you have gotten a broken connection in the middle of a large download, you will appreciate this new feature. After a connection is broken during a download, reconnect to the FTP site and begin downloading the file again in the same local directory. The `ncftp` command resumes downloading

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

where it left off.

Using ncftp for background transfers

If you are moving around to different parts of the FTP site, or jumping between different sites, you might not want to wait around for a file transfer to complete before you can go somewhere else. The `ncftp` command has an excellent feature for placing transfer commands in a spool file and then running them in the background immediately or later.

To select a file for background transfer, use the `bgget` command. Here is an example:

ncftp /pub > **bgget wireless.doc**

- Spooled: get wireless.doc

ncftp /pub > **jobs**

---Scheduled-For-----Host----------Command--------------------

2003-09-01 18:38 maple GET wireless.doc

ncftp /pub > **bgstart**

Background process started.

Watch the "/home/br/.ncftp/spool/log" file to see how it is progressing

In this example, the `bgget` command spools the `wireless.doc` file from the remote current directory (`/pub`) and sets it to be copied to the local current directory. Typing the `jobs` command shows that the job (GET wireless.doc), from the host named `maple`, is scheduled to run immediately (6:38 p.m., September 1, 2003). By running the `bgstart` command this process, and any other spooled jobs, are run immediately.

Instead of starting the background transfers immediately you can do them later. For background jobs spooled for transfer, the transfer begins when you quit `ncftp`, leave your current FTP site or go to another FTP site. In this way, the program that does the transfer (`ncftpbatch`) can take over your current login session to do the transfer.

If you wanted to wait even longer to do the transfers, you can pass an option to the `bgget` command to have it start the transfer at a particular time. Here is an example:

ncftp /pub > **bgget -@ 2003090410000 wireless.doc**

With this command, the transfer is set to run at 1 a.m. on September 4, 2003 to transfer the file called `wireless.doc`. Again, you can check the `.ncftp/spool/log` file in your home directory to see if the transfer has completed.

Using the gFTP window

If you prefer a more graphical interface for accessing FTP servers, you can use the gFTP window. You can open a gFTP window by typing `gftp` or by choosing Internet→ More Internet Applications → gFTP from the Red Hat menu on the desktop. An example of the gFTP window is shown in Figure 9-10.

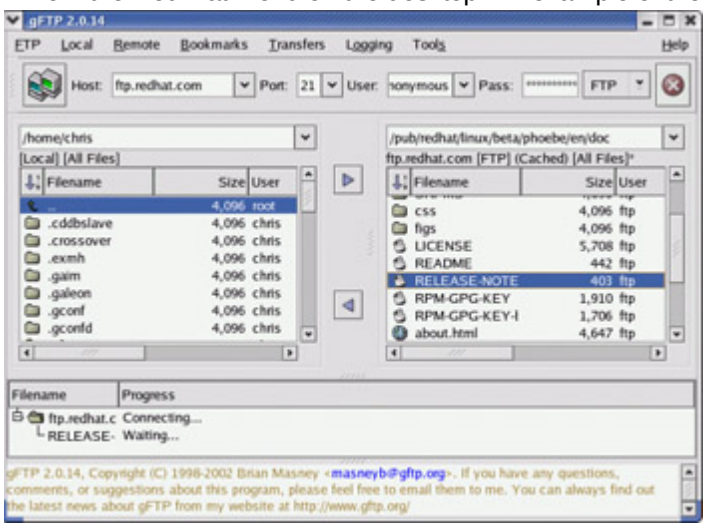

Figure 9-10: View local and remote files simultaneously from the gFTP window.

Unlike the `ftp` command, the gFTP window lets you simultaneously see the contents of the current remote and local directories. To transfer a file from one side to the other, simply double-click it or drag-and-drop it to the other pane. (Normally, you will just be copying files from FTP sites, unless a site

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

provides you with permission to write to it.)

Follow this procedure to connect to an FTP site:

1. Type the name of the FTP server to which you want to connect (for example, `ftp.redhat.com`) into the Host box.
2. Type the port number on the FTP server (you can leave this blank to use the default port number 21).
3. Type the user name used to log in to the FTP server. Use the default `anonymous` if you don't have a specific user name and the server is pubically accessible.
4. Type the password for the user name you entered. The convention with FTP servers is to use your e-mail address as the password.
5. Click the icon displaying two little monitors to connect to the FTP site.
6. If you entered the information correctly, the bottom pane on the window should show that the transfer was complete and the right pane should show the contents of the current directory of the FTP site. Here are some actions you can take once you are connected:

   Move around. Double-click a directory to move to that directory or two dots (. .) to move up a level. You can do this on both the remote and local directories.

   Drag-and-drop files. You can drag-and-drop files from the FTP site on to the left pane (representing your current local directory).

   Save this site. If you want to return to this site later, choose Bookmarks → Add Bookmark. A pop-up window lets you name this site for the bookmarks list. After you do, you can select that entry from the list at a later date to connect to that site. The gFTP window will have stored not only the host name, but also the port, user name, and password. So you are just one click away from connecting.

A nice feature of gFTP is that it stores log information. Choose Logging → View Log. A window appears showing you the conversations that have taken place between your computer and each FTP site. You can look at these messages to see what is wrong if you are unable to connect to a site or to remember where you have been and what you have done on an FTP site.

Getting files with wget

If you already know where a file is on the network, there are more efficient ways of downloading that file than opening an FTP session, moving around the FTP server, and running the `get` command. The `wget` command is a simple, efficient tool for doing non-interactive downloads of files over the Internet.

If there is one or more files you want to download from an FTP site or Web server (HTTP), and you know exactly where those files are, `wget` is a good way to download. The `wget` command is very useful if you want to copy a whole Web site, recursively, from one computer to another. When downloading from FTP sites, `wget` can let you just download as the anonymous user or add your own user name and password to the command line.

Here is an example of using `wget` to get a file from an FTP site:

$ **wget ftp://ftp.redhat.com/pub/contrib/libc6/i386/zsh-4.1.1-1.i386.rpm**

-- 01:53:33- ftp://ftp.redhat.com/pub/contrib/libc6/i386

= `zsch-4.1.1-1.i386.rpm`

Resolving ftp.redhat.com ... done.

Conecting to ftp.redhat.com[216.148.218.201]:21... connected.

Loggin in as anonymous ... Logged in!

==> SYST ... done. ==> PWD ... done.

==> TYPE I ... done. ==> CWD /pub/contrib/libc6/i386 ... done.

==> PASV ... done. ==> RETR zsh-4.1.1-1.i386.rpm... done.

Length: 2,541,655 (unauthoritative)

100%[==========>] 2,541,655 50.31K/s ETA 00:00

05:16:39 (50.31 KB/s) - `zsch-4.1.1-1.i386.rpm` saved [2541655]

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

By the first part of the URL (`ftp://`), `wget` knows you are copying a file from using FTP to the current directory (`.`) on the local host. After resolving the address (`ftp.redhat.com`), `wget` connects to the site, logs in as the anonymous user, changes the working directory (to `/pub/contrib/libc6/i386`), and transfers the file. As the file downloads, `wget` shows the progress of the download, then exits.
If you need to log in as a user other than anonymous, you can add that information to the command line or to a `.netrc` file in your home directory (type **man netrc** to see the format of that file). Here is an example of adding the password to the command line:

$ wget ftp://joe:my67chevy@ftp.handsonhistory.com/memol.doc .

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

Adding a password to a command line leaves the password exposed to onlookers. This practice is generally discouraged, except in cases where no one can see your monitor

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

In the previous example, the user logs in as `joe` with the password of `my67chevy`. The `wget` then copies the file `memo1.doc` from the current directory on the host computer named `ftp.handsonhistory.com`. That current directory is most likely `/home/joe`.

Using `wget`, you can download files from Web servers as well. The `wget` command downloads files using http protocol, if file addresses begin with `http://`. Downloading a single file, you would use the same form as you would for an FTP file (for example: `wget http://host/file.`). The best `wget` option for HTTP downloads is `-r` (recursive).

A recursive download lets you choose a point at a Web site and download all content below that point. Here is an example of a recursive download used to download the contents of the `www.handsonhistory` Web site.

$ **wget –r www.handsonhistory.com .**

In this example, the directory structure making up the `www.handsonhistory.com` Web site is copied below the current directory in the directory `www.handsonhistory.com`. This is useful if you want to move the contents of a Web site from one computer to another.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution

Downloading an entire Web site can result in a massive amount of data being downloaded. If you only want part of a Web site, start from a point lower in the site's structure

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Another command that you might be interested in, similar to wget, is the curl command. Like wget, curl can download files using ftp or http protocols. However, it also includ

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using ssh for remote login/remote execution

Although the `telnet` and `rlogin` login commands and `rsh` remote execution command have been around much longer, the `ssh` command is the preferred tool for remote logins and executions. The reason is that `ssh` provides encrypted communication so you can use it securely over insecure, public networks between hosts that don't know each other.

In the following example, `ssh` is being used to log in to a computer named `maple`. Because no user is specified, `ssh` tries to log in as the current user (which is the root user in this case).

ssh maple

root@maple's password:

If you wanted to log in as a different user, you could use the `-l` option. For example, to log in to the computer named `maple` as the user named `jake`, you could type the following:

ssh -l jake maple

jake@maple's password:

The `ssh` command can also be used to execute a command on the remote computer. For example, if you wanted to monitor the messages file on a remote computer for a minute, you could type the following command:

ssh maple -l root "tail -f /var/log/messages"

root@maple's password:

After you typed the password in the above case, the last several lines of the `/var/log/messages` file on the remote computer would be displayed. As messages were received, they would continue to be displayed until you decided to exit (press Ctrl+D).

Note Find out more about the ssh command from the SSH Web site (www.openssh.org).

Using scp for remote file copy

The `scp` command is a simple yet secure way of copying files among Linux systems. It uses the underlying `ssh` facility, so if `ssh` is enabled, so is `scp`. Here is an example of using `scp` to copy a file from one computer to another:

scp myfile toys.linuxtoys.net:/home/chris

root@toys.linuxtoys.net's password: ******
In this example, the file `myfile` is copied to the computer named `toys.linuxtoys.net` in the `/home/chris` directory. If you don't provide a user name (as is the case here), `scp` assumes you are using the current user name. Unlike some tools that provide remote login, `scp` and `ssh` do allow you to login as root user over the network.
To use `scp` with a different user name, you can append the user name with an `@` character. For example, `chris@toys.linuxtoys.net:/home/chris` would attempt to log in as the user named chris to do the file copy.
The first time you connect to a remote computer using `scp` or `ssh`, those commands try to establish the authenticity of the remote host. If it cannot establish the host's authenticity, it will display the RSA key fingerprint and ask you if you want to continue. If you type `yes`, `scp` will not question the authenticity of that computer again for subsequent `scp` commands.
However, if the RSA key fingerprint should change in the future for the remote computer (which will happen if, for example, the operating system is reinstalled on that computer), `scp` will refuse to let you connect to that remote computer. To override that refusal, you need to edit your `$HOME/.ssh/known_hosts` file and delete the entry for the remote computer. You can then verify the authenticity of the remote computer and continue to use `scp`.

Using the "r" commands: rlogin, rcp, and rsh

The `rlogin`, `rcp`, and `rsh` commands all use the same underlying security mechanism to enable remote login, remote file copy, and remote execution, respectively, among computers. These commands are included with Red Hat Linux to be compatible with legacy UNIX systems. Because "r" commands are inherently insecure, however, most people use `ssh` and `scp` commands to provide the same functionality in a more secure way.

Cross-Reference	For a description of these "r" commands, refer to the *Red Hat Linux Bible* Web site at Wiley Publishing: www.wiley.com/compbooks/negus.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

Most use of the World Wide Web centers on the Web browser. Mozilla is popular with Red Hat Linux, and is accessible from a button on your desktop. The e-mail program that comes with Mozilla, called Mozilla Mail, provides almost everything you need for reading, sending, and managing your mail. Ximian Evolution, which can also be used for managing your e-mail, is now the preferred mail reader for Red Hat Linux. To read newsgroups, you can also use Mozilla or the Pan newsreader.
Commands such as `telnet` (for remote login) and `ftp` (for remote file copying) are available for communicating with remote computers from the command line in Red Hat Linux. There are also improved programs for remote login and file copy, such as `ssh`, `scp`, and `ncftp`.

Team LiB

PREVIOUS NEXT

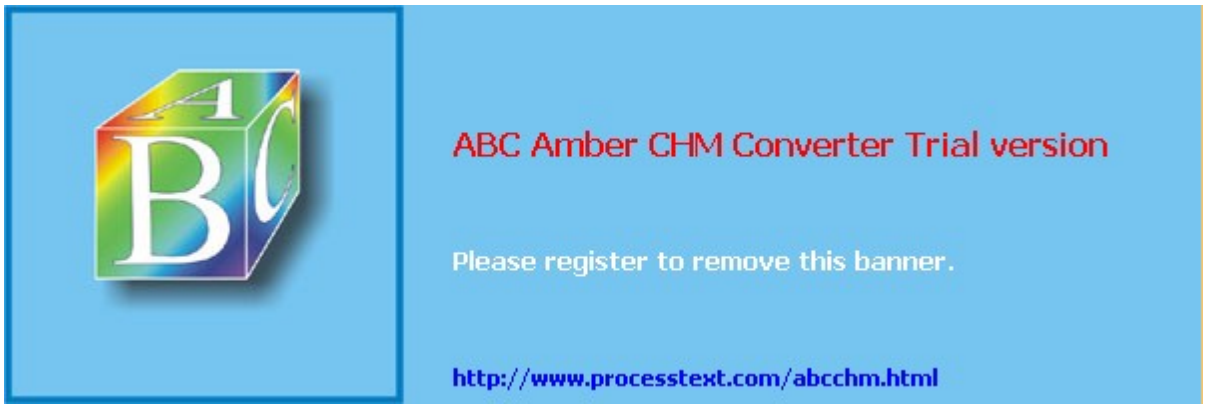

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Part III: Administering Red Hat Linux

Chapter List

Chapter 10: Understanding System Administration
Chapter 11: Setting Up and Supporting Users
Chapter 12: Automating System Tasks
Chapter 13: Backing Up and Restoring Files
Chapter 14: Computer Security Issues

Team LiB PREVIOUS NEXT

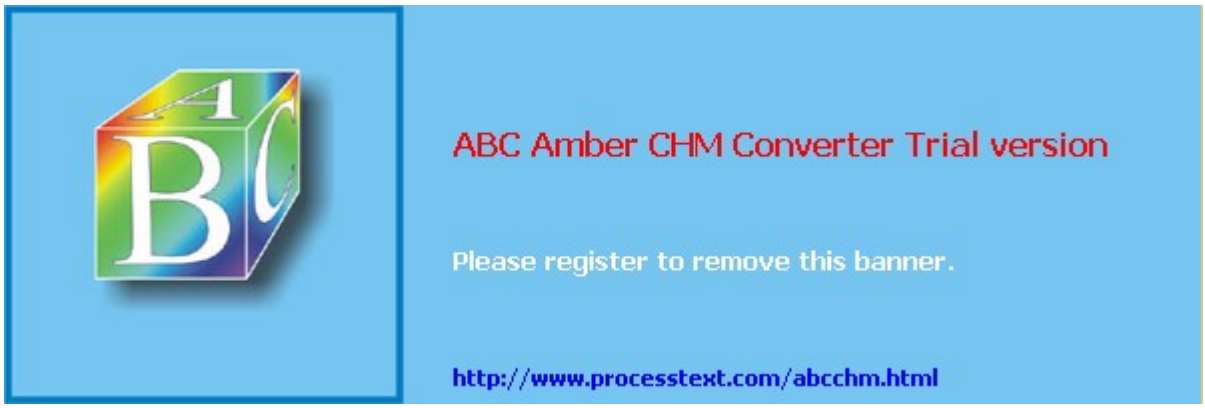

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 10: Understanding System Administration

Overview

Red Hat Linux, like other UNIX systems, was intended for use by more than one person at a time. Multiuser features allow many people to have accounts in Red Hat Linux, with their data kept secure from others. Multitasking allows many people to use the computer at the same time. Sophisticated networking protocols and applications make it possible for a Red Hat Linux computer to extend its capabilities to network users and computers around the world. The person assigned to manage all of this stuff is referred to as the *system administrator.*

Even if you are the only person using a Red Hat Linux system, system administration is still set up to be separate from other computer use. To do most tasks, you need to be logged in as the root user (also referred to as the super user). Other users cannot change, or in some cases, even see some of the configuration information for a Red Hat Linux system. In particular, security features such as passwords are protected from general view.

General principles of Red Hat Linux system administration are described in this chapter. In particular, this chapter covers some of the basic tools you need to administer your Red Hat Linux system. It also helps teach you how to work with file systems and monitor the setup and performance of your Linux system.

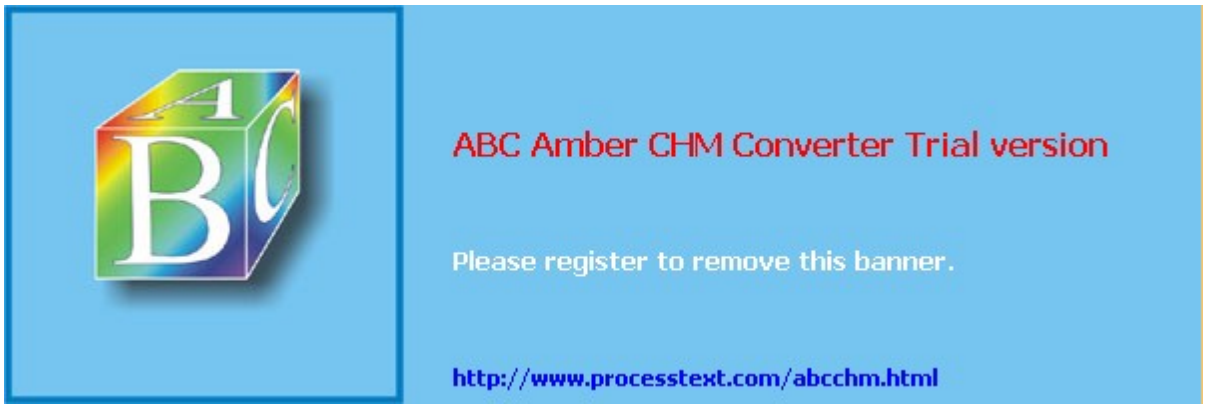

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the root Login

The root user has complete control of the operation of your Red Hat Linux system. That user can open any file or run any program. The root user also installs software packages and adds accounts for other people who use the system.

When you first install Red Hat Linux, you should add a password for the root user. You need to remember and protect this password. You will need it to log in as root or to obtain root permission while you are logged in as some other user.
The home directory for the root user is `/root`. That and other information associated with the root user account is located in the `/etc/passwd` file. Here is what the root entry looks like in the `/etc/passwd` file:

```
root:x:0:0:root:/root:/bin/bash
```

This shows that for the user named `root`, the user ID is set to `0` (root user), the group ID is set to `0` (root group), the home directory is `/root`, and the shell for that user is `/bin/bash`. You can change the home directory or the shell used, if you like, by editing the values in this file.

Cross-Reference See the section on setting up users in Chapter 11 for more information about the /etc/passwd file.

Among the defaults that are set for the root user are aliases for certain dangerous commands. Aliases for the `rm`, `cp`, and `mv` commands allow those commands to be run with the `-i` option. This prevents massive numbers of files being removed, copied, or moved by mistake. The `-i` option causes each deletion, copy, or move to prompt you before the actual change is made.

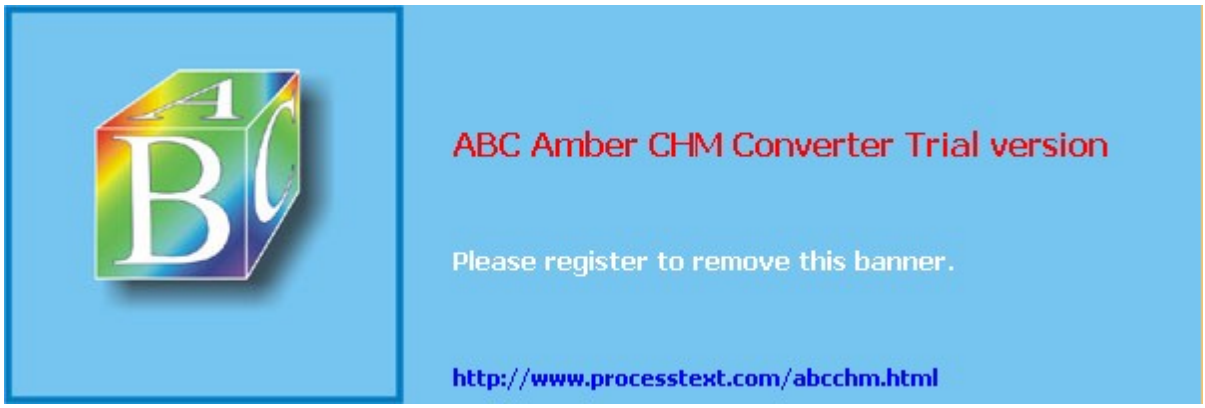

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Becoming Super User (The su Command)

Though the normal way to become the super user is to log in as root, sometimes that is not convenient. For example, you may be logged into a regular user account and just want to make a quick administrative change to your system without having to log out and log back in. Or, you may need to log in over the network to make a change to a Linux system but find that the system doesn't allow root users in from over the network (a common practice).
The answer is that you can use the `su` command. From any Terminal window or shell, you can simply type:

```
$ su
Password: ******
#
```

When you are prompted, type in the root user's password. The prompt for the regular user (`$`) will be changed to the super user prompt (`#`). At this point, you have full permission to run any command and use any file on the system. However, one thing that the `su` command doesn't do when used this way is read in the root user's environment. As a result, you may type a command that you know is available and get the message "command not found." To fix this problem, you can use the `su` command with the dash (-) option instead, as follows:

```
$ su -
Password: ******
#
```

You still need to type the password, but after you do that, everything that normally happens at login for the root user will happen after the `su` command is completed. Your current directory will be root's home directory (probably `/root`), and things like the root user's `PATH` variable will be used. With the other way of running root, you would not have changed directories or the environment of the current login session.

Tip When you become super user during someone else's session, a common mistake is to leave files or directories behind in the user's directories that are owned by root. If you do this, be sure to use the chown or chmod command to make the files and directories you modify open to the user that you want to own them. Otherwise, you will probably get a phone call in a short time, asking you to come back and fix it.

You can also use the `su` command to become another user than root. For example, to have the permissions of a user named `chum`, you could type the following:

```
$ su - chum
```

Even if you were root user before you typed this command, you would only have the permission to open files and run programs that are available to `chum`. As root user, however, after you type the `su` command to become another user, you don't need a password to continue. If you type that command as a regular user, you must type the new user's password.
When you are finished using super user permissions, return to the previous shell by exiting the current shell. Do this by pressing Ctrl+D or by typing **exit**. If you are the administrator for a computer that is accessible to multiple users, don't leave a root shell open on someone else's screen (unless you want to let that person do anything they like to the computer)!
When you run GUI tools as a regular user, you are usually prompted for the root password (as described in an upcoming section, "Using Red Hat graphical administration tools"). If a GUI tool fails and doesn't prompt you for a password, refer to the "Becoming Super User in X" sidebar.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Learning about Administrative GUI Tools, Commands, Configuration Files, and Log Files

Red Hat Linux has advanced enough in recent releases that you can now do most system administration from your desktop GUI, bypassing the shell altogether. Whether you administer Red Hat Linux from the GUI or from a shell, however, underlying your activities are many administrative commands, configuration files, and log files.

[redacted]

Becoming Super User in X

There[redacted]s when an X GUI is running on Red Hat Linux as a non-root user and you want to run a graphical administration program. In most cases, the GUI administration program will simply prompt you for the root password to continue. However, if the program fails, saying that you don't have permission to run the command, here is what you can do:

First, open a Terminal window on the X desktop.

Then, open permission to the X window display to everyone on the local computer (this is just a temporary measure) by typing:

$ **xhost +localhost**

Type the following and enter the root password when prompted, to become super user:

$ **su -**
Password: ******
#

Next, type the following to see the current display value:

echo $DISPLAY
If the value is something like :0 or :0.0, any X command you run from that shell will appear on the console terminal for the computer. If you are at the console and that's what you see, then you can skip the next step. If you see no value (which is quite possible) or the wrong value, you must set the `DISPLAY` variable.
Type the following (assuming you are using a `bash` or `sh` shell):
export DISPLAY=:0
At this point, you can run any administrative X command (such as `neat` or `redhat-config-packages`) and have it appear on your X desktop. If you are running an administrative command from a remote computer and you want it to appear on your local desktop, you can set the `DISPLAY` to *`host`*`:0`, where *`host`* is replaced by the name of your computer.

When you are done, be sure to exit the application you are running. Then restore the security of your X desktop by typing the following:

$ **xhost –**

[redacted]

Understand[redacted]I tools, commands, and files are located and how they are used will help you effectively m[redacted]Red Hat Linux system. Although most administrative features are intended for the root user, other administrative users (described later in this section) have limited administrative

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

capabilities.

Using Red Hat graphical administration tools

The trend over the past few versions of Red Hat Linux has been to steer clear of the massive administrative interfaces (such as `linuxconf` and `Webmin`) and instead to offer graphical windows that perform individual administrative tasks. Instead of sharing one monolithic interface, they share common menus. Individual graphical windows for configuring a network, adding users, or setting up printers can be launched from that menu.

To administer your Red Hat Linux system through the GNOME or KDE desktops, Red Hat has provided a common menu (Red Hat calls it the Main Menu; I refer to it as the Red Hat menu). In GNOME click the red hat; in KDE, click the K icon in the lower-left corner of the desktop.
Because these administrative tasks require root permission, if you are logged in as a regular user you must enter the root password before the GUI application's window opens. For example, if you launch the System Logs window (System Tools → System Logs) from the GNOME menu as a regular user, you see the pop-up window shown in Figure 10-1.

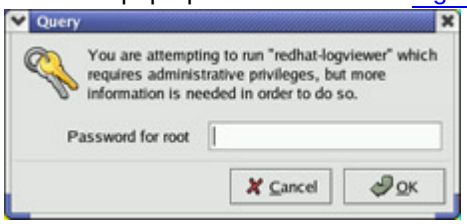

Figure 10-1: Enter the root password to open system administration windows from a regular user's GUI.

After you have entered the root password, most of the Red Hat configuration tools will open without requiring you to retype the password during this login session. Look for a "keys" icon in the lower right corner of the panel, indicating that you have root authorization. Click the keys to open a pop-up window that lets you remove authorization. Otherwise, authorization goes away when you close the GUI window.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note As you configure different features on your Red Hat Linux system, you are asked to launch different individual graphical windows. In general, if you have a choice of tools for configuring

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The following list describes many of the GUI-based windows you can use to administer your Red Hat Linux system. Start these windows from the Server Settings or System Tools submenus on your Red Hat menu:

Server Settings — This submenu accesses the following server configuration windows:

Domain Name Service — Create and configure zones if your computer is acting as a DNS server.

HTTP— Configure your computer as an Apache Web server.

NFS— Set up directories from your system to be shared with other computers on your network using the NFS service.

Samba— Configure Windows (SMB) file sharing. (To configure other Samba features, you can use the SWAT window. From the Red Hat menu, click System Settings→ Server Settings → More Server Settings → Samba Configuration.)

Services — Lets you display and change which services are running on your Red Hat Linux system at different run levels.

Add/Remove Applications — Manage software packages in the Red Hat distribution.

Authentication — Lets you change how users are authenticated on your system. Usually, Shadow Passwords and MD5 Passwords are selected. However, if your network supports LDAP, Kerberos, SMB, NIS, or Hesiod authentication, you can select to use any of those authentication types.

Date & Time — Lets you set the date and time or choose to have an NTP server keep system time in sync. Figure 10-2 shows the Date/Time Properties window.

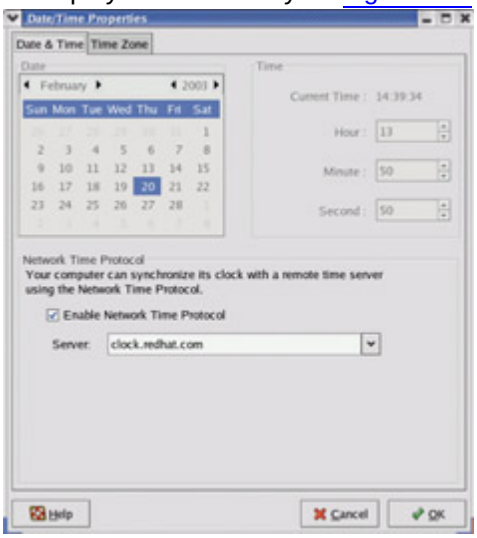

Figure 10-2: Choose an NTP server or set date and time on the Date/Time Properties window.

Disk Management — Mount and format removable media, such as CDs and floppy disks.

Display — Change the settings for your X desktop, including color depth and resolution for your display. You can also choose settings for your video card and monitor.

Hardware Browser — View information about your computer's hardware.

Internet Configuration Wizard — Create initial configurations for connecting to the Internet via Ethernet, ISDN, modem, and other types of network equipment.

Keyboard — Choose the type of keyboard you are using, based on language.

Kickstart — Lets you create a Kickstart configuration file that can be used to install multiple Red Hat Linux systems without user intervention.

Login Screen — Control how your login screen appears and behaves.

Mouse — Configure your mouse.

Network — Lets you manage your current network interfaces, as well as add interfaces.

Network Device Control — Display the active profile for network devices.

Printing Manager — Lets you configure local and network printers.

Red Hat Network — Lets you register your computer with the Red Hat Network to

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

get free software updates.
Root Password — Change the root password.
Security Level — Configure your firewall to allow or deny services to computers from the network.
Soundcard Detection — Tries to detect and configure your sound card.
System Logs — Displays system log files and lets you search them for keywords.
Figure 10-3 shows the Boot Log being displayed in the System Logs window.

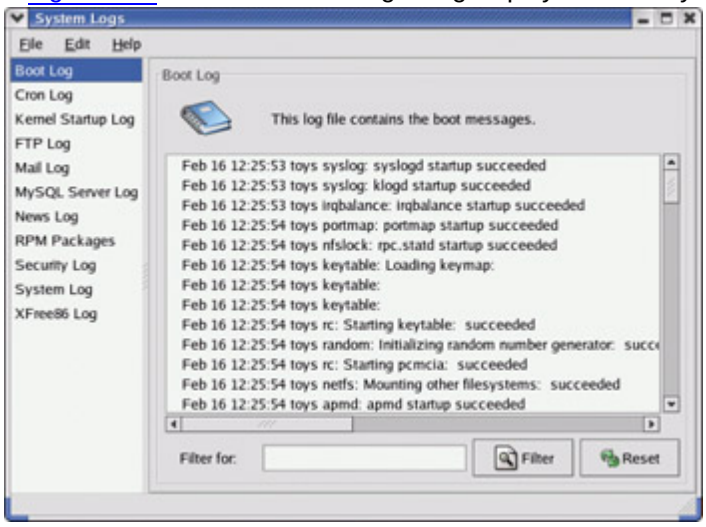

Figure 10-3: Show log files of activities from system boot, FTP, mail, news, and other services.

System Monitor — Shows information about running processes and resource usage.
Task Scheduler — Schedules tasks to be run at set times.
Users & Groups — Lets you add, display, and change user and group accounts for your Red Hat Linux system.

Procedures for using the various Red Hat graphical administrative tools are distributed throughout the book.

Administrative commands

Many commands are intended only for root. When you log in as root, your `$PATH` variable is set to include some directories that contain commands for the root user. These include the following directories:

`/sbin` — This contains commands for modifying your disk partitions (such as `fdisk`), changing boot procedures (`grub`), and changing system states (`init`).
`/usr/sbin` — This contains commands for managing user accounts (such as `useradd`) and configuring your mouse (`mouseconfig`) or system services (`chkconfig`). Commands that run as daemon processes are also contained in this directory. (Look for commands that end in "d" such as `sshd`, `pppd`, and `crond`.)

Some administrative commands are contained in regular user directories (such as `/bin` and `/usr/bin`). This is especially true of commands that have some options available to everyone. An example is the `/bin/mount` command, which anyone can use to list mounted file systems, but only root can use to mount file systems.

To find commands that are intended primarily for the system administrator, check out the section 8 manual pages (usually in `/usr/share/man/man8`). They contain descriptions and options for most Red Hat Linux administrative commands.

Some third-party applications will add administrative commands to directories that are not in your `PATH`. For example, an application may put commands in `/usr/local/bin`, `/opt/bin`, or `/usr/local/sbin`. In those cases, you may want to add those directories to your `PATH`.

Administrative configuration files

Configuration files are another mainstay of Linux administration. Almost everything you set up for your particular computer — user accounts, network addresses, or GUI preferences — is stored in plain-text files. This has some advantages and some disadvantages.

The advantage of plain-text files is that it is easy to read and change them. Any text editor will do. On

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

the downside, however, is that as you edit configuration files, no error checking is going on. You have to run the program that reads these files (such as a network daemon or the X desktop) to find out if you set up the files correctly. A comma or a quote in the wrong place can sometimes cause a whole interface to fail.

Throughout this book, I describe the configuration files you need to set up the different features that make up Red Hat Linux. In terms of a general perspective on configuration files, however, there are several locations in the Red Hat Linux file system where configuration files are stored. Here are some of the major locations:

$HOME — All users store information in their home directories that directs how their login accounts behave. Most configuration files in $HOME begin with a dot (.), so they don't appear as a user's directory when you use a standard `ls` command (you need to type **ls -a** to see them). There are dot files that define how each user's shell behaves, the desktop look and feel , and options used with your text editor. There are even files (such as `.ssh/*` and `.rhosts`) that configure network permissions for each user.

/etc — This directory contains most of the basic Linux system-configuration files. The following `/etc` configuration files are of interest:

adjtime — Holds to data to adjust the hardware clock (see the `hwclock` man page).

aliases — Can contain distribution lists used by the Linux mail service.

bashrc — Sets system-wide defaults for bash shell users. (By default, it sets the shell prompt to include current user name, host name, current directory, and other values.)

cdrecord.conf — Contains defaults used for recording CDs.

crontab — Sets cron environment and times for running automated tasks.

csh.cshrc (or **cshrc**) — Sets system-wide defaults for csh (C shell) users.

exports — Contains a list of local directories that are available to be shared by remote computers using the Network File System (NFS).

fdprm — Sets parameters for common floppy disk formats.

fstab — Identifies the devices for common storage media (hard disk, floppy, CD-ROM, etc.) and locations where they are mounted in the Linux system. This is used by the `mount` command to choose which file systems to mount.

group — Identifies group names and group IDs (GIDs) that are defined on the systems. Group permissions in Red Hat Linux are defined by the second of three sets of rwx (read, write, execute) bits associated with each file and directory.

gshadow — Contains shadow passwords for groups.

host.conf — Sets the locations in which domain names (e.g., redhat.com) are searched for on TCP/IP networks (such as the Internet). By default, the local hosts file is searched, then any nameserver entries in `resolv.conf`.

hosts — Contains IP addresses and host names that you can reach from your computer. (Usually this file is used just to store names of computers on your LAN or small private network.)

hosts.allow — Lists host computers that are allowed to use certain TCP/IP services from the local computer.

hosts.deny — Lists host computers that are *not* allowed to use certain TCP/IP services from the local computer (doesn't exist by default).

inittab — Contains information that defines which programs start and stop when Red Hat Linux boots, shuts down, or goes into different states in between. This is the most basic configuration file for starting Linux.

issue — Contains the lines that are displayed when a terminal is ready to let you log in to Red Hat Linux from a local terminal, or the console in text mode.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

issue.net — Contains login lines that are displayed to users that try to log in to the Linux system from a computer on the network using the telnet service.
lilo.conf — Sets Linux boot loader (lilo) parameters to boot the computer. In particular, it lists information about bootable partitions on your computer. (If you are using grub, which replaced lilo as the default boot manager, the `lilo.conf.anaconda` file is available. You can copy that file to `lilo.conf` to switch to LILO.)
mail.rc — Sets system-wide parameters associated with using mail.
man.config — Used by the `man` command to determine the default path to the location of `man` pages.
modules.conf — Contains aliases and options related to loadable kernel modules used by your computer.
mtab — Contains a list of file systems that are currently mounted.
mtools.conf — Contains settings used by DOS tools in Linux.
named.conf — Contains DNS settings if you are running your own DNS server.
ntp.conf — Includes information needed to run the Network Time Protocol (NTP).
passwd — Stores account information for all valid users for the system. Also includes other information, such as the home directory and default shell.
printcap — Contains definitions for the printers configured for your computer.
profile — Sets system-wide environment and start-up programs for all users. This file is read when the user logs in.
protocols — Sets protocol numbers and names for a variety of Internet services.
redhat-release — Contains a string identifying the current Red Hat release.
resolv.conf — Identifies the locations of DNS name server computers that are used by TCP/IP to translate Internet host.domain names into IP addresses.
rpc — Defines remote procedure call names and numbers.
services — Defines TCP/IP services and their port assignments.
shadow — Contains encrypted passwords for users that are defined in the `passwd` file. (This is viewed as a more secure way to store passwords than the original encrypted password in the `passwd` file. The `passwd` file needs to be publicly readable, whereas the `shadow` file can be unreadable by all but the root user.)
shells — Lists the shell command-line interpreters (`bash`, `sh`, `csh`, etc.) that are available on the system, as well as their locations.
sudoers — Sets commands that can be run by users, who may not otherwise have permission to run the command, using the `sudo` command. In particular, this file is used to provide selected users with root permission.
syslog.conf — Defines what logging messages are gathered by the syslogd daemon and what files they are stored in. (Typically, log messages are stored in files contained in the `/var/log` directory.)
termcap — Lists definitions for character terminals, so that character-based applications know what features are supported by a given terminal. Graphical terminals and applications have made this file obsolete to most people. (Termcap was the BSD UNIX way of storing terminal information; UNIX System V used definitions in `/usr/share/terminfo` files.)
xinetd.conf — Contains simple configuration information used by the `xinetd` daemon process. This file mostly points to the `/etc/xinetd.d` directory for information about individual services (described later).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

/etc/X11 — Contains subdirectories that each contain system-wide configuration files used by X and different X window managers available with Red Hat Linux. The `XF86Config` file (which makes your computer and monitor usable with X) and configuration directories containing files used by `xdm` and `xinit` to start X are in here.

Directories relating to window managers contain files that include the default values that a user will get if that user starts one of these window managers on your system. Window managers that may have system-wide configuration files in these directories include GNOME (`gdm`) and Twm (`twm`).

Note Some files and directories in /etc/X11 are linked to locations in the /usr/X11R6 directory.

/etc/alternatives — Contains links that the alternatives facility uses to enable a system administrator to exchange one service with another in a way that is invisible to users. (Currently, only mail and printing use the alternatives service.)

/etc/amanda — Contains files and directories that allow the amanda facility to do network backups of other Linux and UNIX systems.

/etc/cipe — Holds if-up and if-down scripts to start a CIPE virtual private network.

/etc/cron* — Directories in this set contain files that define how the `crond` utility runs applications on a daily (`cron.daily`), hourly (`cron.hourly`), monthly (`cron.monthly`), or weekly (`cron.weekly`) schedule.

/etc/cups — Contains files that are used to configure the CUPS printing service.

/etc/default — Contains files that set default values for various utilities. For example, the file for the `useradd` command defines the default group number, home directory, password expiration date, shell, and skeleton directory (`/etc/skel`) that are used when creating a new user account.

/etc/httpd — Contains a variety of files used to configure the behavior of your Apache Web server (specifically, the `httpd` daemon process).

/etc/init.d — Contains the permanent copies of run-level scripts. These scripts are linked to files in the `/etc/rc?.d` directories to have each service associated with a script started or stopped for the particular run level. The *?* is replaced by the run-level

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

number (0 through 6).
/etc/mail — Contains files used to configure your sendmail mail service.
/etc/pcmcia — Contains configuration files that allow you to have a variety of PCMCIA cards configured for your computer. (PCMCIA slots are those openings on your laptop that allow you to have credit-card-sized cards attached to your computer. You can attach such devices as modems and external CD-ROMs.)
/etc/postfix — Contains configuration files for the postfix mail transport agent.
/etc/ppp — Contains several configuration files used to set up Point-to-Point protocol (so that you can have your computer dial out to the Internet).
/etc/rc?.d — There is a separate `rc?.d` directory for each valid system state: `rc0.d` (shutdown state), `rc1.d` (single-user state), `rc2.d` (multiuser state), `rc3.d` (multiuser plus networking state), `rc4.d` (user-defined state), `rc5.d` (multiuser, networking, plus GUI login state), and `rc6.d` (reboot state).
/etc/security — Contains files that set a variety of default security conditions for your computer. These files are part of the pam (pluggable authentication modules) package.
/etc/skel — Any files contained in this directory are automatically copied to a user's home directory when that user is added to the system. By default, most of these files are dot (.) files, such as .kde (a directory for setting KDE desktop defaults) and `.bashrc` (for setting default values used with the bash shell).
/etc/squid — Contains configuration files for the Squid proxy caching server.
/etc/sysconfig — Contains important system configuration files that are created and maintained by various Red Hat applications (including iptables, `samba`, and most networking services).
/etc/uucp — Contains configuration files used with Taylor UUCP (a nonstandard version of the uucp facility that is used to create modem, direct line, and other serial connections with other computers).
/etc/xinetd.d — Contains a set of files, each of which defines a network service that the `xinetd` daemon listens for on a particular port. When the `xinetd` daemon process receives a request for a service, it uses the information in these files to determine which daemon processes to start to handle the request.

Administrative log files

One of the things that Linux does well is keep track of itself. This is a good thing, when you consider how much can go wrong with a complex operating system. Sometimes you are trying to get a new facility to work and it fails without giving you the foggiest reason why. Other times you want to monitor your system to see if people are trying to access your computer illegally. In any of those cases, you can use log files to help track down the problem.
The main utilities for logging error and debugging messages for Linux are the `syslogd` and `klogd` daemons. General system logging is done by `syslogd`. Logging that is specific to kernel activity is done by `klogd`. Logging is done according to information in the `/etc/syslog.conf` file. Messages are typically directed to log files that are usually in the `/var/log` directory.
Red Hat Linux includes a System Logs window (System Tools→ System Logs) you can use to view and search system log files from the desktop. See Chapter 14 for a description of that window and of the log files you can view with it.

Using other administrative logins

You don't hear much about other administrative logins (besides root) being used with Red Hat Linux. It was a fairly common practice in UNIX systems to have several different administrative logins that allowed administrative tasks to be split among several users. For example, a person sitting near a printer could have lp permissions to move print jobs to another printer if they knew a printer wasn't working.

In any case, these administrative logins are available with Linux, so you may want to look into using them. At the very least, because individual software packages such as bind, squid, and amanda set up permissions for their log files and configuration files based on their administrative logins, maintaining those permissions can prevent someone who hacks into one of those services from gaining control of the whole computer.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Because most Red Hat Linux administrative features are expected to be administered by the root user, e-mail for other administrative accounts is routed to the root user. If you want

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Understanding administrative logins

Here are some of the administrative logins that are configured automatically for Linux systems. By tradition, these logins are assigned UID numbers under 100. Here are examples:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Most administrative logins have no passwords by default. They also typically have /sbin/nologin assigned as their shell, so if you try to login as one of these users, you see a "Th

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

lp — This user can control some printing features. Having a separate lp administrator allows someone other than the super user to do such things as move or remove lp logs and print spool files. The home directory for lp is `/var/spool/lpd`.
mail — This user can work with administrative e-mail features. The mail group has group permissions to use mail files in `/var/spool/mail` (which is also the mail user's home directory).
uucp — This user owns various `uucp` commands (once used as the primary method for dial-up serial communications). It is the owner of log files in `/var/log/uucp`, spool files in `/var/spool`, administrative commands (such as `uuchk`, `uucico`, `uuconv`, and `uuxqt`) in `/usr/sbin`, and user commands (`uucp`, `cu`, `uuname`, `uustat`, and `uux`) in `/usr/bin`. The home directory for uucp is `/var/spool/uucp`.
bin — This user owns many commands in `/bin` in traditional UNIX systems. This is not the case in Red Hat Linux, because root tends to own most executable files. The home directory of bin is `/bin`.
news — This user could do administration of Internet news services, depending on how you set permission for `/var/spool/news` and other news-related resources. The home directory for news is `/var/spool/news`.

Using sudo for assigning administrative privilege

One way to give full or limited root privileges to any non-root user is to set up the `sudo` facility. That simply entails adding the user to `/etc/sudoers` and defining what privilege you want that user to have. Then the user can run any command he or she is privileged to use by preceding that command with the `sudo` command.
The following is an example of how to use the `sudo` facility to cause any users that are added to the `wheel` group to have full root privileges:

1. As the root user, edit the `/etc/sudoers` file by running the `visudo` command:
2. # **visudo**
   By default, the file is opened in `vi`, unless your EDITOR variable happens to be set to some other editor acceptable to `visudo` (for example, `export EDITOR=gedit`) The reason for using `visudo` is that the command will lock the `/etc/sudoers` file and do some basic sanity-checking of the file to ensure it was edited correctly.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you are stuck here, refer to the vi tutorial in Chapter 4 for information on using the vi editor.

3. Uncomment the following line to allow users in the group named `wheel` to have full root privileges on the computer:

   ```
   %wheel  ALL=(ALL)  ALL
   ```

   The previous line causes the user to be prompted for a password to be allowed to use administrative commands. To allow users in the wheel group to have that privilege without using a password, uncomment the following line instead:

   ```
   %wheel  ALL=(ALL)  NOPASSWD: ALL
   ```

4. Save the changes to the `/etc/sudoers` file (in `vi`, type **ZZ**).
5. Still as root user, open the `/etc/group` file in any text editor and add the users you want to have root privilege to the `wheel` line. For example, if you were to add the users `mary` and `jake` to the `wheel` group, the line would appear as follows:

   ```
   wheel:x:10:root,mary,jake
   ```

At this point, the users mary and jake can run the `sudo` command to run commands, or parts of commands, that are normally restricted to the root user. The following is an example of a session by the user jake after he has been assigned `sudo` privileges:

```
[jake]$ sudo umount /mnt/win

We trust you have received the usual lecture from the local System
Administrator. It usually boils down to these two things:

    #1) Respect the privacy of others.
    #2) Think before you type.
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Password: *********

[jake]$ **mount /mnt/win**

mount: only root can mount /dev/hda1 on /mnt/win

[jake]$ **sudo mount /mnt/win**

[jake]$

In the above session, the user `jake` runs the `sudo` command so he can unmount the `/mnt/win` file system (using the `umount` command). He is given a warning and asked to provide his password (this is jake's password, *not* the root password).

Notice that even after `jake` has given the password, he must still use the `sudo` command to run the command as root (the first mount fails, but the second succeeds). Notice that he was not prompted for a password for the second `sudo`. That's because after entering his password successfully he can enter as many `sudo` commands as he wants for the next five minutes without having to enter it again. (You can change the timeout value from five minutes to however long you want by setting the `passwd_timeout` value in the `/etc/sudoers` file.)

The preceding example grants a simple all-or-nothing administrative privilege to everyone you put in the `wheel` group. However, the `/etc/sudoers` file gives you an incredible amount of flexibility in permitting individual users and groups to use individual applications or groups of applications. I recommend you refer to the `sudoers` and `sudo` man pages for information about how to tune your `sudo` facility.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Administering Your Red Hat Linux System

Your Linux system administrator duties don't end after you have installed Red Hat Linux. Your ongoing job as a Linux system administrator includes the following tasks:

Configuring hardware — Often when you add hardware to your Red Hat Linux computer, that hardware will be automatically detected and configured by tools such as kudzu. In those cases where the hardware was not properly set up, you can use commands such as `lsmod`, `insmod`, and `rmmod` to configure the right modules to get the hardware working.

Managing file systems and disk space — You must keep track of the disk space being consumed, especially if your Red Hat Linux system is shared by multiple users. At some point, you may need to add a hard disk or track down what is eating up your disk space (you can use commands like `find` to do this).

Monitoring system performance — You may have a run-away process on your system or you may just be experiencing slow performance. Tools that come with Red Hat Linux can help you determine how much of your CPU and memory are being consumed.

Keeping software up2date — Corrections to Red Hat Linux software, especially those related to security issues, should be incorporated into your system as time goes on. The Red Hat Network offers the up2date service to enusre that you get critical fixes.

The aforementioned administrative tasks are described in the rest of this chapter. Later chapters cover other administrative topics, such as managing user accounts (Chapter 11), automating system tasks (Chapter 12), system backups and restores (Chapter 13), and securing your system (Chapter 14). Tasks related to network administration are covered in Chapters 15 through 26.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

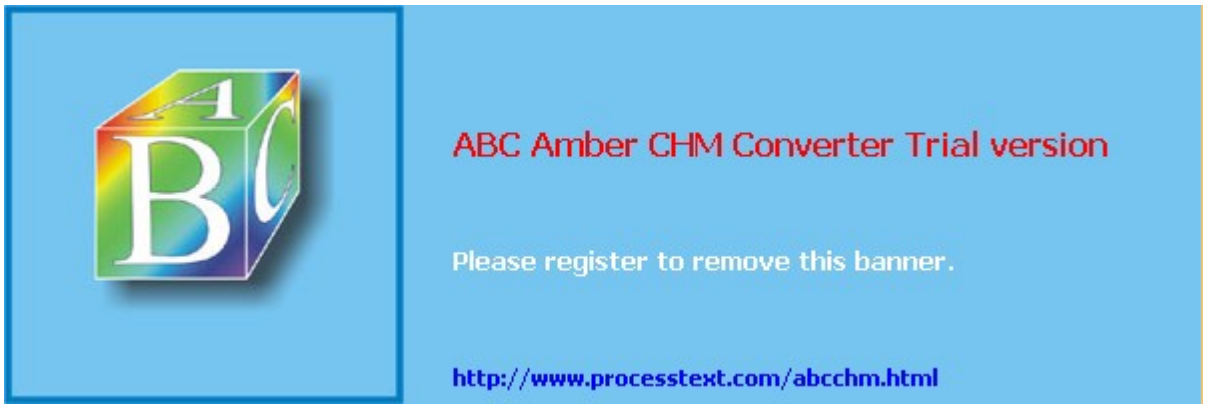

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Managing File Systems and Disk Space

File systems in Red Hat Linux are organized in a hierarchy, beginning from root (/) and continuing downward in a structure of directories and subdirectories. As an administrator of a Red Hat Linux system, it is your duty to make sure that all the disk drives that represent your file system are available to the users of the computer. It is also your job to make sure there is enough disk space in the right places in the file system for users to store what they need.

File systems are organized differently in Linux than they are in MS Windows operating systems. Instead of drive letters (for example, A:, B:, C:) for each local disk, network file system, CD-ROM, or other type of storage medium, everything fits neatly into the directory structure. It is up to an administrator to create a mount point in the file system and then connect the disk to that point in the file system.

Cross-Reference

Chapter 2 provides instructions for using Disk Druid to configure disk partitions. Chapter 4 describes how the Linux file system is organized.

The organization of your file system begins when you install Linux. Part of the installation process is to divide your hard disk (or disks) into partitions. Those partitions can then be assigned to:

- A part of the Linux file system,
- Swap space for Linux, or
- Other file system types (perhaps containing other bootable operating systems)

For our purposes, I want to focus on partitions that are used for the Linux file system. To see what partitions are currently set up on your hard disk, use the `fdisk` command as follows:

fdisk –l

Disk /dev/hda: 40.0 GB, 40020664320

255 heads, 63 sectors/track, 4825 cylinders

Units = cylinders of 16065 * 512 bytes = 8225280 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/hda1	*	1	13	104	b	Win95 FAT32
/dev/hda2		84	89	48195	83	Linux
/dev/hda3		90	522	3478072+	83	Linux
/dev/hda4		523	554	257040	5	Extended
/dev/hda5		523	554	257008+	82	Linux swap

This output shows the disk partitioning for a computer able to run both Red Hat Linux and Microsoft Windows. You can see that the Linux partition on `/dev/hda3` has most of the space available for data. There is a Windows partition (`/dev/hda1`) and a Linux swap partition (`/dev/hda5`). There is also a small /boot partition (46MB) on `/dev/hda2`. In this case, the root partition for Linux has 3.3GB of disk space and resides on `/dev/hda3`.

Next, to see what partitions are actually being used for your Linux system, you can use the `mount` command (with no options). The `mount` command can show you which of the available disk partitions are actually mounted and where they are mounted.

mount

/dev/hda3 on / type ext3 (rw)

/dev/hda2 on /boot type ext3 (rw)

/dev/hda1 on /mnt/win type vfat (rw)

none on /proc type proc (rw)

none on /dev/pts type devpts (rw,gid=5,mode=620)

/dev/cdrom on /mnt/cdrom type iso9660 (ro,nosuid,nodev)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note You may notice that /proc, /dev/pts, and other entries not relating to a partition are shown as file systems. This is because they represent different file system types (proc and devpts, respectively). The word none, however, indicates that they are not associated with a separate partition.

The mounted Linux partitions in this case are `/dev/hda2`, which provides space for the `/boot` directory (which contains data for booting Linux), and `/dev/hda3`, which provides space for the rest of the Linux file system beginning from the root directory (`/`). This particular system also contains a Windows partition that was mounted in the `/mnt/win` directory and a CD that was mounted in its standard place: `/mnt/cdrom`. (With most GUI interfaces, the CD is typically mounted automatically when you insert it.)

After the word `type`, you can see the type of file system contained on the device. (See the description of different file system types later in this chapter.) Particularly on larger Linux systems, you may have multiple partitions for several reasons:

Multiple hard disks — You may have several hard disks available to your users. In that case you would have to mount each disk (and possibly several partitions from each disk) in different locations in your file system.

Protecting different parts of the file system — If you have many users on a system, and the users consume all of the file system space, the entire system can fail. For example, there may be no place for temporary files to be copied (so the programs writing to temporary files may fail), and incoming mail may fail to be written to mail boxes. With multiple mounted partitions, if one partition runs out, others can continue to work.

Backups — There are some fast ways to back up data from your computer that involve copying the entire image of a disk or partition. If you want to restore that partition later, you can simply copy it back (bit-by-bit) to a hard disk. With smaller partitions, this approach can be done fairly efficiently.

Protecting from disk failure — If one disk (or part of one disk) fails, by having multiple partitions mounted on your file system, you may be able to continue working and just fix the one disk that fails.

When a disk partition is mounted on the file system, all directories and subdirectories below that mount point are then stored on that partition. So, for example, if you were to mount one partition on `/` and one on `/usr`, everything below the `/usr` mount point would be stored on the second partition while everything else would be stored on the first partition. If you then mounted another partition on `/usr/local`, everything below that mount point would be on the third partition, while everything else below `/usr` would be on the second partition.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip

What if a remote file system is unmounted from your computer, and you go to save a file in that mount point directory? What happens is that you will write the file to that dir

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Mount points that are often mentioned as being candidates for separate partitions include `/`, `/boot`, `/home`, `/usr`, and `/var`. The root file system (`/`) is the catchall for directories that aren't in other mount points. The root file system's mount point (/) is the only one that is required. The `/boot` directory holds the images needed to boot the operating system. The `/home` file systems is where all the user accounts are typically stored. Applications and documentation are stored in `/usr`. Below the `/var` mount point is where log files, temporary files, server files (Web, FTP, and so on), and lock files are stored (that is, items that need disk space for your computer's applications to keep running).

Cross-Reference See Chapter 2 for further information on partitioning techniques.

The fact that multiple partitions are mounted on your file system is invisible to people using your Linux system. The only times they will care will be if a partition runs out of space or if they need to save or use information from a particular device (such as a floppy disk or remote file system). Of course, any user can check this by typing the `mount` command.

Mounting file systems

Most of your hard disks are mounted automatically for you. When you installed Red Hat Linux, you were asked to create partitions and indicate the mount points for those partitions. When you boot Red Hat Linux, all Linux partitions should be mounted. For that reason, this section focuses mostly on how to mount other types of devices so that they become part of your Red Hat Linux file system.
Besides being able to mount other types of devices, you can also use `mount` to mount other kinds of file systems on your Linux system. This means that you can store files from other operating systems or use file systems that are appropriate for certain kinds of activities (such as writing large block sizes). The most common use of this feature for the average Linux user, however, is to allow that user to obtain and work with files from floppy disks or CD-ROMs.

Supported file systems

To see file system types that are currently in use on your system, type **cat /proc/filesystems**. The following file system types are supported in Linux:

befs — This is the file system used by the BeOS operating system.
ext3 — The ext file systems are the most common file systems used with Linux. The ext3 file system was new for Red Hat Linux 7.2 and is the default file system type. The root file system (`/`) must be ext3, ext2, or minux. The ext3 file system is also referred to as the Third Extended file system. The ext3 file system includes journaling features that improve a file system's ability to recover from crashes, as compared to ext2 file systems.
ext2 — The default file system type for previous versions of Red Hat Linux. Features are the same as ext3, except that ext2 doesn't include journaling features.
ext — This is the first version of ext3. It is not used very often anymore.
iso9660 — This file system evolved from the High Sierra file system (which was the original standard used on CD-ROM). Extensions to the High Sierra standard (called Rock Ridge extensions), allow iso9660 file systems to support long filenames and UNIX-style information (such as file permissions, ownership, and links).
kafs — This is the AFS client file system. It is used in distributed computing environments to share files with Linux, Windows, and Macintosh clients.
minix — This is the Minix file system type, used originally with the Minix version of UNIX. It only supports filenames of up to 30 characters.
msdos — This is an MS-DOS file system. You can use this type to mount floppy disks that come from Microsoft operating systems.
umsdos — This is an MS-DOS file system with extensions to allow features that are similar to UNIX (including long filenames).
proc — This is not a real file system, but rather a file-system interface to the Linux kernel. You probably won't do anything special to set up a proc file system. However, the `/proc` mount point should be a proc file system. Many utilities rely on `/proc` to gain access to Linux kernel information.
reiserfs — This is the ReiserFS journaled file system.
swap — This is used for swap partitions. Swap areas are used to temporarily hold data when RAM is currently used up. Data is swapped to the swap area, then

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

returned to RAM when it is needed again.
nfs — This is the Network File System (NFS) type of file system. File systems mounted from another computer on your network use this type of file system.

Cross-Reference Information on using NFS to export and share file systems over a network is contained in Chapter 18.

hpfs — This file system is used to do read-only mounts of an OS/2 HPFS file system.
ncpfs — This relates to Novell NetWare file systems. NetWare file systems can be mounted over a network.

Cross-Reference For information on using NetWare file systems over a network, see the section on setting up a file server in Chapter 18.

ntfs — This is the Windows NT file system. It is supported as a read-only file system (so that you can mount and copy files from it). Read-write support is available, but not built into the kernel by default and is considered unreliable (some say, dangerous).
affs — This file system is used with Amiga computers.
ufs — This file system is popular on Sun Microsystems operating systems (i.e., Solaris and SunOS).
xenix — This was added to be compatible with Xenix file systems (one of the first PC versions of UNIX). The system is obsolete and will probably be removed eventually.
xiafs — This file system supports long filenames and larger inodes than file systems such as minux.
coherent — This is the file system type used with Coherent or System V files. Like the xenix file system type, it will be removed at some time in the future.

Using the fstab file to define mountable file systems

The hard disks on your local computer and the remote file systems you use every day are probably set up to automatically mount when you boot Linux. The definitions for which of these file systems are mounted are contained in the `/etc/fstab` file. Here's an example of an `/etc/fstab` file:

```
LABEL=/      /           ext3    defaults            1 1
LABEL=/boot  /boot       ext3    defaults            1 2
none         /dev/pts    devpts  gid=5,mode=620      0 0
/dev/fd0     /mnt/floppy auto    noauto,owner        0 0
none         /proc       proc    defaults            0 0
/dev/hda5    swap        swap    defaults            0 0
/dev/cdrom   /mnt/cdrom  iso9660 noauto,owner,kudzu,ro 0 0
/dev/hda1    /mnt/win    vfat    noauto              0 0
```

All file systems listed in this file are mounted at boot time, except for those set to `noauto` in the fourth field. In this example, the root (/) and boot (/boot) hard disk partitions are mounted at boot time, along with the `/proc` and `/dev/pts` file systems (which are not associated with particular devices). The floppy disk (`/dev/fd0`) and CD-ROM drives (`/dev/cdrom`) are not mounted at boot time. Definitions are put in the fstab file for floppy and CD-ROM drives so that they can be mounted in the future (as described later).
I also added one additional line for `/dev/hda1`, which allows me to mount the Windows (vfat) partition on my computer so I don't have to always boot Windows to get at the files on my Windows partition.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note To access the Windows partition described above, I must first create the mountpoint (by typing **mkdir /mnt/win**). I can then mount it when I choose by typing (as root) **mount /mnt/win**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You find the following in each field of the fstab file:

Field 1 — The name of the device representing the file system. The word `none` is often placed in this field for file systems (such as `/proc` and `/dev/pts`) that are not associated with special devices. Notice that this field can now include the `LABEL` option. Using `LABEL`, you can indicate a universally unique identifier (UUID) or volume label instead of a device name. The advantage to this approach is that, since the partition is identified by volume name, you can move a volume to a different device name and not have to change the `fstab` file.

Field 2 — The mount point in the file system. The file system contains all data from the mount point down the directory tree structure, unless another file system is mounted at some point beneath it.

Field 3 — The file system type. Valid file system types are described in the "Supported file systems" section earlier in this chapter.

Field 4 — Options to the `mount` command. In the preceding example, the `noauto` option prevents the indicated file system from being mounted at boot time. Also, `ro` says to mount the file system read-only (which is reasonable for a CD-ROM player). Commas must separate options. See the `mount` command manual page (under the `-o` option) for information on other supported options.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Normally, only the root user is allowed to mount a file system using the mount command. However, to allow any user to mount a file system (such as a file system on a floppy disk), you can

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Field 5 — The number in this field indicates whether or not the indicated file system needs to be dumped. A number 1 assumes that the file system needs to be dumped. A number 2 assumes that the file system doesn't need to be dumped.

Field 6 — The number in this field indicates whether or not the indicated file system needs to be checked with `fsck`. A number 1 assumes that the file system needs to be checked. A number 2 assumes that the file system doesn't need to be checked.

If you want to add an additional local disk or an additional partition, you can create an entry for the disk or partition in the `/etc/fstab` file. To get instructions on how to add entries for an NFS file system, see Chapter 18.

Using the mount command to mount file systems

Your Red Hat Linux system automatically runs `mount -a` (mount all file systems) each time you boot. For that reason, you would typically only use the `mount` command for special situations. In particular, the average user or administrator uses `mount` in two ways:

To display the disks, partitions, and remote file systems that are currently mounted.

To temporarily mount a file system.

Any user can type the `mount` command (with no options) to see what file systems are currently mounted on the local Linux system. The following is an example of the `mount` command. It shows a single hard disk partition (`/dev/hda1`) containing the root (`/`) file system, and proc and devpts file system types mounted on `/proc` and `/dev`, respectively. The last entry shows a floppy disk, formatted with a standard Linux file system (ext3) mounted on the `/mnt/floppy` directory.

$ **mount**

/dev/hda3 on / type ext3 (rw)

none on /proc type proc (rw)

/dev/hda2 on /boot type ext3 (rw)

none on /dev/pts type devpts (rw,gid=5,mode=0620)

/dev/fd0 on /mnt/floppy type ext3 (rw)

The most common devices to mount by hand are your floppy disk and your CD-ROM. However, depending on the type of desktop you are using, CD-ROMs and floppy disks may be mounted for you automatically when you insert them. (In some cases, the autorun program may also run automatically. For example, autorun may start a CD music player or software package installer to handle the data on the medium.)

If you want to mount a file system manually, however, the `/etc/fstab` file helps make it simple to mount a floppy disk or a CD-ROM. In some cases, you can use the `mount` command with a single option to indicate what you want to mount, and information is taken from the `/etc/fstab` file to fill in the other options. Entries probably already in your `/etc/fstab` file let you do these quick mounts in the following two cases:

CD-ROM — If you are mounting a CD-ROM that is in the standard ISO 9960 format (as most software CD-ROMs are), you can mount that CD-ROM by placing it in your CD-ROM drive and typing the following:

mount /mnt/cdrom

By default, your CD-ROM is mounted on the `/mnt/cdrom` directory. (The file system type, device name, and other options are filled in automatically.) To see the contents, type **`cd /mnt/cdrom`**, then type **`ls`**. Files from the CD-ROM's root directory will be displayed.

Floppy Disk — If you are mounting a floppy disk that is in the standard Linux file system format (ext3), you can mount that floppy disk by inserting it in your floppy drive and typing the following:

mount /mnt/floppy

The file system type (ext3), device (`/dev/fd0`), and mount options are filled in from the `/etc/fstab` file. You should be able to change to the floppy disk directory (`cd /mnt/floppy`) and list the contents of the floppy's top directory (`ls`).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note In both of the two previous cases, you could give the device name (/dev/cdrom or /dev/fd0, respectively) instead of the mount point directory to get the same results.

Of course, it is possible that you may get floppy disks you want to use that are in all formats. Someone may give you a floppy containing files from a Microsoft operating system (in MS-DOS format). Or you may get a file from another UNIX system. In those cases, you can fill in your own options, instead of

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

relying on options from the `/etc/fstab` file. In some cases, Linux autodetects that the floppy disk contains an MS-DOS (or Windows vfat) file system and mount it properly without additional arguments. However, if it doesn't, here is an example of how to mount a floppy containing MS-DOS files:

mount -t msdos /dev/fd0 /mnt/floppy

This shows the basic format of the `mount` command you would use to mount a floppy disk. You could change the msdos to any other supported file system type (described earlier in this chapter) to mount a floppy of that type. Instead of using floppy drive A: (`/dev/fd0`), you could use drive B: (`/dev/fd1`) or any other accessible drive. Instead of mounting on `/mnt/floppy`, you could create any other directory and mount the floppy there.

Here are some other useful options you could add along with the `mount` command:

-t auto — If you aren't sure exactly what type of file system is contained on the floppy disk (or other medium you are mounting), use the `-t auto` option to indicate the file system type. The `mount` command will query the disk to try to guess what type of data it contains.

-r — If you don't want to make changes to the mounted file system (or can't because it is a read-only medium), use this option when you mount it. This will mount it read-only.

-w — This mounts the file system with read/write permission.

Some options to mount are available only for a specific file system type. See the `mount` manual page for those and other useful options.

Using the umount command to unmount a file system

When you are done using a temporary file system, or you want to unmount a permanent file system temporarily, you can use the `umount` command. This command detaches the file system from its mount point in your Red Hat Linux file system. To use `umount`, you can give it either a directory name or a device name. For example:

umount /mnt/floppy

This unmounts the device (probably `/dev/fd0`) from the mount point `/mnt/floppy`. You could also have done this using the form:

umount /dev/fd0

In general, it's better to use the directory name,.because the `umount` will fail if the device is mounted in more than one location.

If you get a message that the "device is busy," the umount request has failed. The reason is that either a process has a file open on the device or that a you have a shell open with a directory on the device as a current directory. Stop the processes or change to a directory outside of the device you are trying to unmount for the umount request to succeed.

An alternative for unmounting a busy device is the `-l` option. With `umount -l` (a lazy unmount), the unmount happens as soon as the device is no longer busy. To unmount a remote NFS file system that is no longer available (for example, the server went down), you can use the `umount -f` option to forcibly unmount the NFS file system.

Using the mkfs command to create a file system

It is possible to create a file system, for any supported file system type, on a disk or partition that you choose. This is done with the `mkfs` command. While this is most useful for creating file systems on hard disk partitions, you can create file systems on floppy disks or re-writable CDs as well.

Here is an example of using `mkfs` to create a file system on a floppy disk:

mkfs -t ext3 /dev/fd0

mke2fs 1.34, (25-Jul-2003)

Filesystem label=

OS type: Linux

Block size=1024 (log=0)

Fragment size=1024 (log=0)

184 inodes, 1440 blocks

72 blocks (5.00%) reserved for the super user

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

First data block=1

1 block group

8192 blocks per group, 8192 fragments per group

184 inodes per group

Writing inode tables: done

Filesystem too small for a jounal

Writing superblocks and filesystem accounting information: done

The filesystem will be automatically checked every 32 mounts or

180 days, whichever comes first. Use tune2fs -c or -i to override.

You can see the statistics that are output with the formatting done by the `mkfs` command. The number of inodes and block created are output. Likewise, the number of blocks per group and fragments per group are also output. You could now mount this file system (`mount /mnt/floppy`), change to it as your current directory (`cd /mnt/floppy`), and create files on it as you please.

Adding a hard disk

Adding a new hard disk to your computer so that it can be used by Linux requires a combination of steps described in previous sections. The general steps are as follows:

1. Install the hard disk hardware.
2. Identify the partitions on the new hard disk.
3. Create the file systems on the new hard disk.
4. Mount the file systems.

The easiest way to add a hard disk to Linux is to have the entire hard disk devoted to a single Linux partition. You can have multiple partitions, however, and assign them each to different types of file systems and different mount points, if you like. The procedure below describes how to add a hard disk containing a single Linux partition. Along the way, however, it also notes which steps you need to repeat to have multiple file systems with multiple mount points.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note This procedure assumes that Red Hat Linux is already installed and working on the computer. If this is not the case, follow the instructions for adding a hard disk on your current

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

1. Install the hard disk into your computer. Follow the manufacturer's instructions for physically installing and connecting the new hard disk. If, presumably, this is a second hard disk, you may need to change jumpers on the hard disk unit itself to have it operate as a slave hard disk. You may also need to change the BIOS settings.
2. Boot your computer to Red Hat Linux.
3. Determine the device name for the hard disk. As root user from a shell, type:
4. # **dmesg | less**

   From the output, look for an indication that the new hard disk was found. For example, if it is a second IDE hard disk, you should see `hdb:` in the output. For a second SCSI drive, you should see `sdb:` instead. Be sure you identify the right disk or you will erase all the data from disks you probably want to keep!
5. Use the `fdisk` command to create partitions on the new disk. For example, if you are formatting the second IDE disk (hdb), you could type the following:
6. # **fdisk /dev/hdb1**
7. If the disk had existing partitions on it, you can change or delete those partitions now. Or, you can simply reformat the whole disk to blow everything away. Use `p` to view all partitions and `d` to delete a partition.
8. To create a new partition, type the following:
9.
10. **n**

    You are asked to choose an extended or primary partition.
11. To choose a primary partition, type the following:
12.
13. **p**

    You are asked the partition number.
14. If you are creating the first partition (or for only one partition), type the number one:
15.
16. **1**

    You are asked to enter the first cylinder number (with one being the default).
17. To begin at the second cylinder, type the number two as follows:
18.
19. **2**

    You are asked to enter the last cylinder.
20. If you are using the entire hard disk, use the last cylinder number shown. Otherwise, choose the ending cylinder number or indicate how many MB the partition should have.
21. To create more partitions on the hard disk, repeat the previous five steps for each partition.
22. Type **w** to write changes to the hard disk. At this point, you should be back at the shell.
23. To make a file system on the new disk partition, use the `mkfs` command. By default, this command creates an ext2 file system, which is useable by Linux. To create an ext2 file system on the first partition of the second hard disk, type the following:
24. # **mkfs -t ext3 /dev/hdb1**

    If you created multiple partitions, repeat this step for each partition (such as `/dev/hdb2`, `/dev/hdb3`, and so on).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you don't use -text3 as shown above, an ext2 file system is created by default. Use other commands, or options to this command, to create other file system types. For example

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

25. Once the file system is created, you can have the partition permanently mounted by editing the `/etc/fstab` and adding the new partition. Here is an example of a line you might add to that file:
26. /dev/hdb1 /abc ext3 defaults 1 1
    In this example, the partition (`/dev/hdb1`) is mounted on the `/abc` directory as an ext3 file system. The `defaults` keyword causes the partition to be mounted at boot time. The numbers `1 1` cause the disk to be checked for errors. Add one line like the one shown above for each partition you created.
27. Create the mount point. For example, to mount the partition on `/abc` (as shown in the previous step), type the following:
28. # **mkdir /abc**
    Create your other mount points if you created multiple partitions. The next time you boot Red Hat Linux, the partition will be automatically mounted on the `/abc` directory, as will any other partitions you added.

Once you have created the file systems on your partitions, a nice tool for adjusting those file systems is the `tune2fs` command. Using `tune2fs`, you can change volume labels, how often the file system is checked, and error behavior. You can also use `tune2fs` to change an ext2 file system to an ext3 file system so the file system can use journaling. For example:

tune2fs -j /dev/hdb1

tune2fs 1.34, (25-Jul-2003)

Creating journal inode: done

This filesystem will be automatically checked every 38 mounts or

180 days, whichever comes first. Use tune2fs -c or -i to override.

By adding the `-j` option to `tune2fs`, you can change either the journal size or attach the file system to an external journal block device. After you have used `tune2fs` to change your file system type, you probably need to correct your `/etc/fstab` file to include changing the file system type from ext2 to ext3.

Using RAID disks

RAID (Redundant Arrays of Independent Disks) is used to spread the data used on a computer across multiple disks, while appearing to the operating system as if it is dealing with a single disk partition. Using the different RAID specifications, you can achieve the following advantages:

Improved disk performance — RAID0 uses a feature called *striping*, where data is striped across multiple RAID partitions. Striping can improve disk performance by spreading the hits on a computer's file system across multiple partitions, which are presumably on multiple hard disks.

Mirroring — RAID1 uses partitions from multiple hard disks as mirrors. That way, if one of the partitions becomes corrupted or the hard disk goes down, the data exists on a partition from another disk because it has continuously maintained an exact mirror image of the original partition.

Parity — Although striping can improve performance, it can increase the chance of data loss, since any hard-disk crash in the array can potentially cause the entire RAID device to fail. Using a feature called *parity*, information about the layout of the striped data is kept so that data can be reconstructed if one of the disks in the array crashes. RAID3, 4, and 5 implement different levels of parity.

During installation of Red Hat Linux, you can use the Disk Druid window to create RAID0, RAID1, and RAID5 disk arrays. The following procedures describe how to set up RAID disks during installation.

Before you begin creating RAID partitions when you install Red Hat Linux, you will probably want to start with a computer that has two or more hard disks. The reason is that, if you don't have multiple hard disks, you won't get the performance gains that come from spreading the hits on your computer among multiple disks. Likewise, mirroring will be ineffective if all RAID partitions are on the same disk, because the failure of a single hard disk would still potentially cause the mirrored partitions to fail as well.

For example, you might begin with 30GB of free disk space on your first hard disk (`/dev/hda`) and 30GB of free disk space on your second hard disk (`/dev/hdb`). During the Red Hat Linux installation procedure, select to partition your disk with Disk Druid. Then follow the procedure that follows:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

1. From the Disk Setup (Disk Druid) window, click the RAID button. A RAID Options window appears.
2. Select Create a Software RAID Partition, and click OK. The Add Partition window appears.
3. With Software RAID selected as the File System Type, choose the drive you want to create the partition on and choose the size (in megabytes). Then click OK.
4. Repeat steps 2 and 3 (presumably creating software RAID partitions on different hard disks until you have created all the RAID partitions you want to use).
5. Click the RAID button again. Select Create a RAID Device [default=/dev/md0] and click OK. The Make RAID Device window appears, as shown in Figure 10-4.

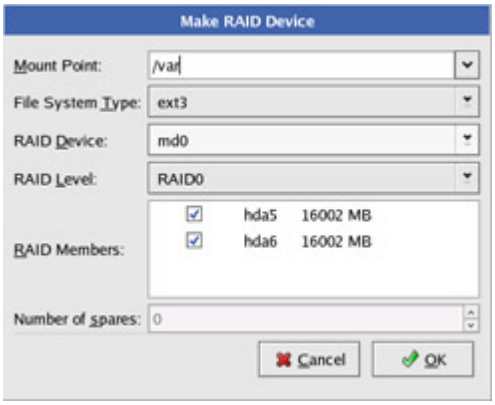

Figure 10-4: Join multiple RAID partitions to form a single RAID device.

6. You need to select the following information about your RAID device and click OK:

   Mount Point — The point in the file system associated with the RAID device. You might be creating the RAID device for a part of the file system for which you expect there to be a lot of hits on the hard disk (such as the `/var` partition).
   File System Type — For a regular Linux partition, choose ext3. You can also select LVM, swap, or VFAT as the file system type.
   RAID Device — The first RAID device is typically md0 (for `/dev/md0`).
   RAID Level — Allowable RAID levels are RAID0, RAID1, and RAID5. RAID0 is for striping (essentially dividing the RAID device into stripes across RAID partitions you have selected). RAID1 is for mirroring so that the data is duplicated across all RAID partitions. RAID5 is for parity, so there is always one backup disk if a disk goes bad. You need at least three RAID partitions to use RAID5.
   RAID Members — From the RAID partitions you created, select which ones are going to be members of the RAID device you are creating.
   Number of spares — Select how many spares are available to the RAID device.

The new RAID Device should appear on the Disk Setup (Disk Druid) window. Click Next to continue with installation.

Once installation is complete, you can check your RAID devices by using a variety of tools that come with Red Hat Linux. Commands for working with RAID partitions come in the raidtools package. Using raidtools commands, you can list and reconfigure your RAID partitions.

Checking system space

Running out of disk space on your computer is not a happy situation. Using tools that come with Red Hat Linux, you can keep track of how much disk space has been used on your computer, and you can keep an eye on users who consume a lot of disk space.

Displaying system space with df

You can display the space available in your file systems using the `df` command. To see the amount of space available on all of the mounted file systems on your Linux computer, type **df** with no options:

$ **df**

```
Filesystem  1k-blocks     Used  Available  Use%   Mounted on
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
/dev/hda3    30645460  2958356    26130408   11%    /
/dev/hda2       46668     8340       35919   19%    /boot
/dev/fd0         1412       13        1327    1%    /mnt/floppy
```

The output here shows the space available on the hard disk partition mounted on the root partition (`/dev/hda1`), /boot partition (`/dev/hda2`), and the floppy disk mounted on the `/mnt/floppy` directory (`/dev/fd0`). Disk space is shown in 1K blocks. To produce output in a more human-readable form, use the `-h` option as follows:

```
$ df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/hda3             29G  2.9G   24G  11% /
/dev/hda2             46M  8.2M   25M  19% /boot
/dev/fd0             1.4M   13k  1.2M   1% /mnt/floppy
```

With the `df -h` option, output appears in a friendlier megabyte or gigabyte listing. Other options with `df` let you:

- Print only file systems of a particular type (`-t type`)
- Exclude file systems of a particular type (`-x type`)
- Include file systems that have no space, such as `/proc` and `/dev/pts` (`-a`)
- List only available and used inodes (`-i`)
- Display disk space in certain block sizes (`--block-size=#`)

Checking disk usage with du

To find out how much space is being consumed by a particular directory (and its subdirectories), you can use the `du` command. With no options, `du` lists all directories below the current directory, along with the space consumed by each directory. At the end, `du` produces total disk space used within that directory structure.

The `du` command is a good way to check how much space is being used by a particular user (`du /home/user1`) or in a particular file system partition (`du /var`). By default, disk space is displayed in 1K block sizes. To make the output more friendly (in kilobytes, megabytes, and gigabytes), use the `-h` option as follows:

```
$ du -h /home/jake
114k    /home/jake/httpd/stuff
234k    /home/jake/httpd
137k    /home/jake/uucp/data
701k    /home/jake/uucp
1.0M    /home/jake
```

The output shows the disk space used in each directory under the home directory of the user named jake (`/home/jake`). Disk space consumed is shown in kilobytes (`k`) and megabytes (`M`). The total space consumed by `/home/jake` is shown on the last line.

Finding disk consumption with find

The `find` command is a great way to find file consumption of your hard disk using a variety of criteria. You can get a good idea of where disk space can be recovered by finding files that are over a certain size or were created by a particular person.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note You must be root user to run this command effectively, unless you are just checking your personal files.

In the following example, the `find` command searches the root file system (`/`) for any files owned by the user named jake (`-user jake`) and prints the filenames. The output of the `find` command is then listed with a long listing in size order (`ls -ldS`). Finally that output is sent to the file `/tmp/jake`. When you read the file `/tmp/jake`, you will find all of the files that are owned by the user `jake`, listed in size order. Here is the command line:

find / -user jake -print -xdev | xargs ls -ldS > /tmp/jake

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip The -xdev option prevents file systems other than the selected file system from being searched. This is a good way to cut out a lot of junk that may be output from the /proc file system. It

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The next example is similar to the previous one, except that instead of looking for a user's files, this command line looks for files that are larger than 100 kilobytes (`-size 100k`):

find / -size 100k -print -xdev | xargs ls -ldS > /tmp/size

You can save yourself a lot of disk space by just removing some of the largest files that are no longer needed. Open the `/tmp/size` file in this example and large files are sorted by size.

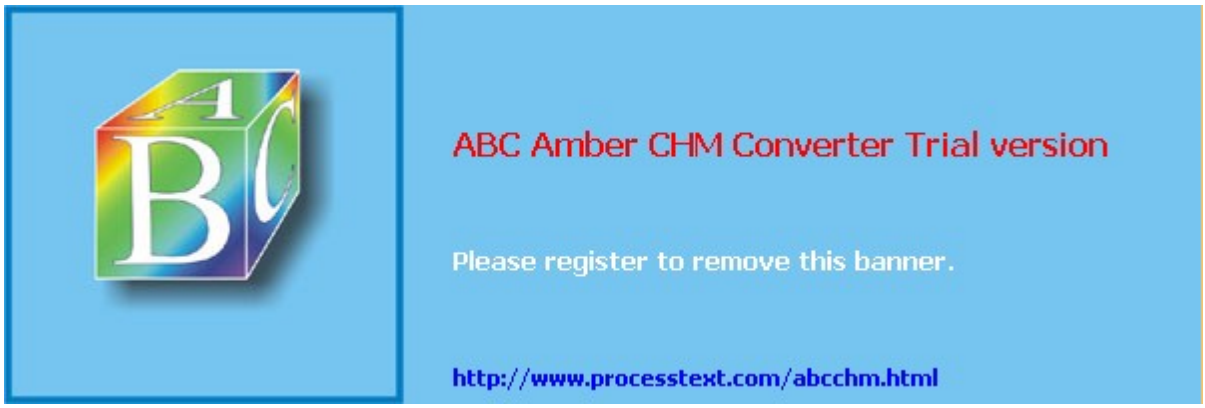

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Monitoring System Performance

If your Linux system is being used as a multiuser computer, sharing the processing power of that computer can be a major issue. Likewise, any time you can stop a runaway process or reduce the overhead of an unnecessary program running, your Linux server can do a better job serving files, Web pages, or e-mail to the people that rely on it.

Utilities are included with Linux that can help you monitor the performance of your Linux system. The kinds of features you want to monitor in Linux include CPU usage, memory usage (RAM and swap space), and overall load on the system. The following sections describe tools for monitoring Linux.

Watch computer usage with System Monitor

If you like visual representations of your system use, the System Monitor provides a great way to see how much your system is being used. To open the System Monitor from the Red Hat menu, click System Tools → System Monitor. Figure 10-5 shows the System Monitor window.

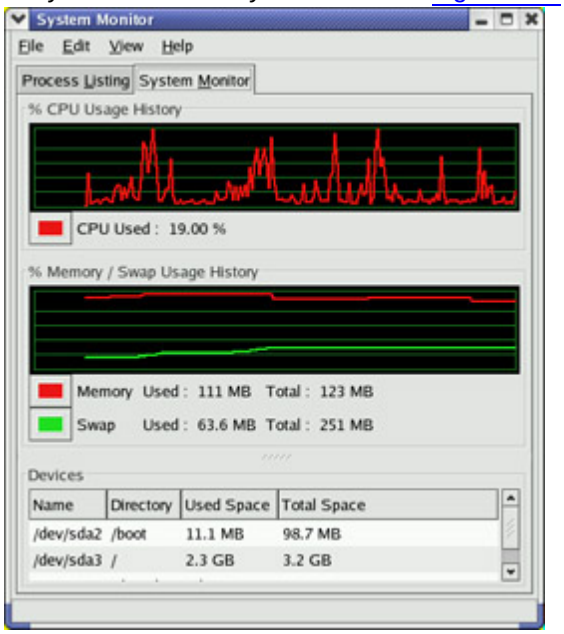

Figure 10-5: System Monitor graphically displays your system's CPU and memory usage.

On the System Monitor window, lines scroll from right to left, indicating the percentage of your CPU being used as it rises and falls. You can also see how much of your total memory (RAM) is being used at the moment (and over time), as well as the amount of swap space being used. Scroll the devices at the bottom of the window to see the space being consumed on each of your hard disk partitions.

Click the Process Listing tab to see a listing of processes that are running currently. Click the columns in that tab to sort processes by name, user (who launched the process), memory use, percentage of CPU being consumed by the process, and process ID.

The example in Figure 10-5 shows a computer that is running with only 128 MB of RAM. The Memory Used is running near the maximum, indicating the likely occurrence of performance problems as RAM fills up and data has to be moved to swap space. The CPU used is only about 19 percent, which indicates that the CPU is keeping up with the demand.

Monitoring CPU usage with top

Start the `top` utility in a Terminal window, and it displays the top CPU consuming processes on your computer. Every five seconds, `top` will determine which processes are consuming the most CPU time and display them in descending order on your screen.

By adding the `-S` option to `top`, you can have the display show you the cumulative CPU time that the process, as well as any child processes that may already have exited, has spent. If you want to change how often the screen is updated, you can add the `-d secs` option, where `secs` is replaced by the number of seconds between updates.

By default, processes are sorted by CPU usage. You can sort processes numerically by PID (press N), by age (press A), by resident memory usage (press M), by time (press T), or back to CPU usage (press

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

P). Figure 10-6 shows an example of `top` running in a Terminal window.

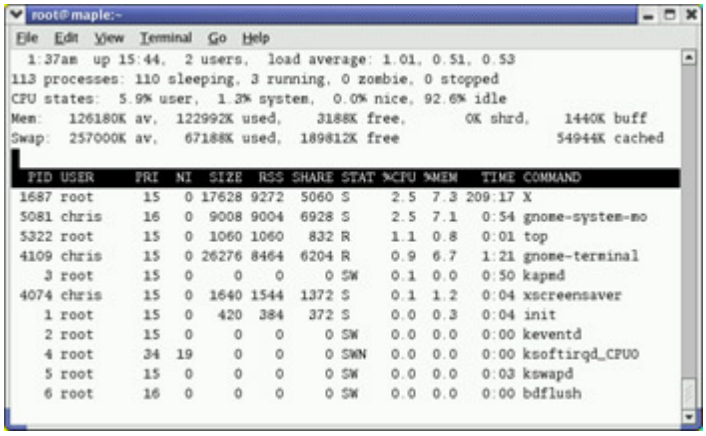

Figure 10-6: Running processes appear in CPU usage order by default in the top window.

Monitoring power usage on laptop computers

To effectively use a laptop computer, you need to be able to monitor and manage the laptop's power usage. Using tools provided in Red Hat Linux, you can configure your laptop to:

- Monitor the battery level
- Notify you when the battery is low
- Notify you when the battery is fully charged
- Show when the laptop is plugged in
- Suspend the current session

Red Hat Linux offers two facilities that do power management: APM and ACPI.

- **Advanced Power Management (APM)** — APM can be used to monitor the battery of your notebook and notify user-level programs to tell you when your battery is low. It can also be used to place your laptop into suspend mode.
- **Advanced Configuration and Power Interface (ACPI)** — Besides monitoring power features on your laptop, ACPI can also do thermal control, motherboard configuration, and change power states.

Many older laptops do not include support for ACPI in the BIOS, so you must use APM to monitor and manage your batteries. For some newer laptops, ACPI may be required. In general, ACPI offers a more complete feature set for power management, but APM has more user-level support today.
To check whether ACPI or APM are supported on your Red Hat Linux system, you can use the `dmesg` command after a reboot. For example, type:

```
# dmesg | less
```

Page through the output looking for lines beginning with ACPI or amp. On a computer where APM wasn't working, I saw the message "amp: BIOS not found." When ACPI wasn't working, I saw the message "ACPI: System description tables not found."
The following procedure was performed on a laptop that used APM to manage power events. It describes how to use the Battery Charge Monitor applet on the desktop to monitor your battery and the `apm` command to place your laptop in suspend mode.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If it seems that ACPI is interfering with the proper operation of your laptop, you can turn off ACPI when you boot your computer. Add acpi=off to the end of the kernel line (from

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the battery status applet

If you are using the GNOME desktop, you can add a Battery Status Monitor to your panel to keep track of the power levels of your battery. The following procedure steps you through adding the monitor to your panel and configuring it to behave as you like:

1. Right-click the GNOME panel; then select Add to Panel → Utility → Battery Charge Monitor. If the Battery Charge Monitor applet appears on the panel (which it should), you can stop here if you like. The applet will:
   Include a battery icon that appears green (with the battery above a 40% charge), yellow (25%-40%), orange (15%-25%) or red (under 15%).
   Show a power cord if the laptop is plugged in.
   Notify you with at pop-up window when the battery charge is low.
   Notify you with a pop-up window when the battery is fully charged.
   Display the percent charged and whether the battery is currently charging when you move the mouse pointer over the applet icon.
   Continue with the following steps if you want to change that behavior.
2. Right-click the Battery Charge monitor applet, then select Preferences. Figure 10-7 shows the Battery Charge Monitor Preferences window and the applet in the panel.

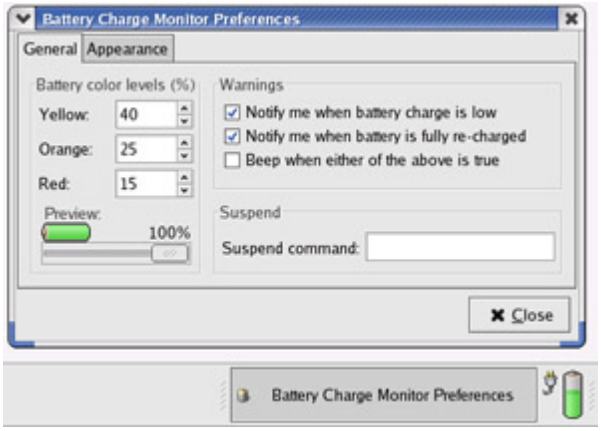

Figure 10-7: View battery status and change preferences with the Battery Charge Monitor.

3. Change any of the following values related to your battery monitor:
   Battery color levels — You can change the levels at which the battery icon changes from green, yellow, orange, and red.
   Warnings — Select whether or not you want to be notified (with a pop-up) when the battery charge is low or fully charged. You can also select to have the laptop beep if either case is true.
   Suspend command — Add a command to put the laptop in suspend mode. (Try the `apm -s` command as described in the next section.) With a Suspend command set, you can right-click the icon and choose Suspend Computer to suspend your laptop.
   Appearance — Select the Appearance tab. From there, you can choose which icons appear on the applet (battery, power cord, and/or percent charged). You can also change whether the progress bar on the battery moves up or down as power empties.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If, when you try to add the Battery Charge Monitor icon to the panel, a pop-up window tells you that the operation fails, it may be that the apmddaemon isn't running. You can

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using apm to enter suspend mode

The `apm` command lets you view information about your computer's power management and put the computer in suspend mode (if it is supported on your laptop). Here are some examples of using the `apm` command:

apm -m

Using the `-m` option, the `apm` command displays the number of minutes of battery life remaining (if that information is available). It may also give you information about the status of how the battery is charging and whether or not the laptop is currently plugged in.

apm -s

The `-s` option of `apm` causes the laptop to enter suspend mode. You can start up the laptop again, in most cases, by pressing a key on the keyboard.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Choosing Software Alternatives

Because several software packages are available for every major service available in Linux (such as mail, printing, and so on), there will be times when people using the system will prefer one service over another. Software packages that have been designed to work with the alternatives system can be configured in Red Hat Linux to let an administrator choose which of the alternatives to a particular service he or she wants to use by default.

Selecting mail and printing alternatives

Beginning with Red Hat Linux 7.3, which was the first version to offer the alternatives feature, two major services were configured to use alternatives: mail transport and printing services. The alternatives facility lets users choose the following, related to mail transport and printing:

Mail Transport Agent (MTA) — If both the sendmail and postfix mail transport agents are installed, as an administrator you can choose which of those two services is the default for sending e-mail.

Printing — If both LPRng and CUPS printing services are installed, you can choose which service is the default for printing documents.

As an administrator, you still need to configure each alternative service to work. Descriptions for configuring sendmail and postfix mail-transport agents are contained in Chapter 19. Information on setting up the CUPS printing service is in Chapter 17. (LPRng is no longer delivered with Red Hat Linux, although it is still available from sites such as `rpmfind.net`.)

In terms of setting up the alternatives side of mail and printing services, much of the work of creating links so that the services can be chosen has already been done. Links relating to the default services are set up in the `/etc/alternatives` directory. Definitions that identify the alternative components of mail and printing are contained in the `/var/lib/alternatives` directory.

Because much of the configuration has been done in advance, the first step in switching between the different printing or mail services installed on your computer is only a couple of clicks away. To switch the default mail or printing services on your computer, do the following:

1. Select one of the following from the Red Hat menu, depending on whether you want to switch your mail or printing service:
   To switch print service, select System Settings → More System Settings → Printer System Switcher.
   To switch mail service, select System Tools → More System Tools → Mail Transport Agent Switcher.
   The switcher window for the service you requested appears.
2. Click on the service you want to switch to — LPRng or CUPS for printing, Sendmail or Postfix for mail. (If the one you want is already selected, you can just cancel.)
   If the switch is successful, a pop-up window tells you to restart the new service.
3. Close the pop-up window.

The next time your computer boots, your new printing or mail service takes over that service. All the links are in place and the start-up scripts are changed. However, your system is still running the old service. The start-up scripts for those four services are in the `/etc/init.d` directory. They are as follows:

`lpd`	For the LPRng printing service
`cups`	For the CUPS printing service
`sendmail`	For the Sendmail service
`postfix`	For the Postfix service

To stop the old service so that the new one can take over, type the following (replacing *service* with the name of the service you want to stop):

/etc/init.d/*service* stop

To start the new service, type the following (replacing *service* with the name of the service you want to start):

/etc/init.d/*service* start

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Providing that the new service was configured properly, it should now be available to the users of your computer.

Using mail and printing alternatives

The printing and mail-transport services that the alternatives facility allows you to change rely on many of the same command names. For example, both printing services use the `lpr` command to print and the `lpq` command to see a printer queue. Both Postfix and Sendmail have a `newalias` command and `mailq` commands for updating aliases and checking the mail queue, respectively.

So, to the user, a change in mail or printing services should (in theory) be nearly invisible. Users can print documents using the same commands they always used and the documents will be printed. Users can send mail as they always did and the fact that a different mail transport is being used should make no difference.

Team LiB

PREVIOUS NEXT

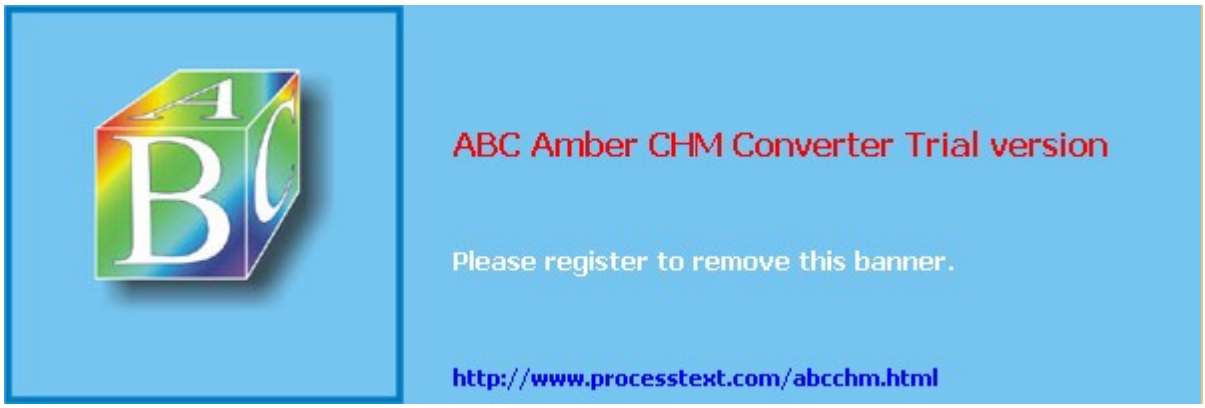

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Getting Linux Software "up2date"

The up2date facility is the preferred method from Red Hat, Inc. for updating Fedora and Red Hat Enterprise Linux software. Using up2date, you can determine if any software packages in your Linux distribution have updates available, and you can ask to have them downloaded and installed automatically.

Fixes and improvements continue after a distribution, such as Fedora Core or Red Hat Enterprise Linux, is released. You can incorporate those fixes and features into Red Hat Linux, using a connection to the Internet and `up2date`.

This section tells you how to determine which updates are available to your Red Hat Linux system. Then it describes how you can register to use the Red Hat Network and use the `up2date` tool to gather those updates.

Note The up2date service is used for getting fixes to your Red Hat Linux system. It is not used to pick up every change that occurs to the software packages included with Red Hat Linux.

With the transition from Red Hat Linux to Fedora and Red Hat Enterprise Linux, changes in what up2date supports are among the most significant changes to the latest distributions from Red Hat. In particular, up2date has been modified to allow you to not only download software from Red Hat FTP sites, but also from repositories of software built to work with the `apt` and `yum` facilities.

Cross-Reference For detailed descriptions of apt and yum, refer to Chapter 5.

Getting updates from the Red Hat Network

Red Hat has created a system that automatically notifies you of available updates and lets you easily download and install them on your Red Hat Linux system. The service is called Red Hat Network. Registering with the Red Hat Network allows you to run the `up2date` program from your Red Hat Linux computer whenever you are ready to do the updates.

Red Hat Network provides a mechanism that can:

- Alert you to the latest fixes and enhancements available for Red Hat Linux.
- Install the packages that include those changes to your system.

A lot of information about the Red Hat Network is available from its Web page (`rhn.redhat.com`). This section provides an overview of how the service works and steps you through the registration process.

Note If you are using an older Red Hat Linux version, you may need to get a new SSL certificate and new up2date software to get the up2date facility to work. Refer to rhn.redhat.com to get those items.

If you don't mind giving out your e-mail address, street address, and phone number, there are some pretty good reasons for creating a Red Hat Network (RHN) account for yourself. With a free RHN account you can register your Red Hat Linux system and create system profiles. Just registering your first machine also entitles you to receive all updates for that machine. To get more services (or entitle more machines), you can purchase a subscription to RHN. To keep your RHN account, Red Hat is currently asking that you occasionally fill out a survey, describing how you use Red Hat Linux.

Note Because Red Hat Network is fairly new and constantly evolving, it is possible that many of the features and entitlements described in this section will change. I recommend that you check the Red Hat Network site for information on new features if the procedures in this section seem outdated.

Using the Red Hat Network alert notification tool

To register your system with Red Hat and start the process of getting notifications of updates, you don't have to go further than your desktop. A round icon with an exclamation point in the middle appears on the desktop panel, ready to help you check for critical updates to Red Hat Linux.

Note If the Red Hat Network Alert icon is not on your desktop, you can put it there by selecting the Red Hat menu on your desktop panel, then clicking System Tools → Red Hat Network Alert Icon.

Figure 10-8 shows you what the Red Hat Network Alert icon looks like when updates are available. (If no

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

updates are available, a blue checkmark icon appears instead.)

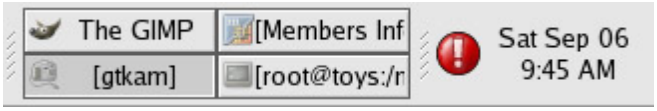

Figure 10-8: The Red Hat Network alert notification tool appears as a round icon on your desktop panel.

Click the icon and the Red Hat Network Alert Notification Tool window appears. Follow these steps to begin the process of getting updates:

1. From the Red Hat Network Alert Notification Tool window, click Forward. The Terms of Service window appears.
2. Read the Terms of Service. Then either click Forward (to continue) or click Remove From Panel (to remove the icon and not use the service). If you continue, the Proxy Configuration window appears.
3. If you need to go through a proxy server to get to the Internet for HTTP (Web) service, click Enable HTTP Proxy, identify the location of the proxy server, and then add any authentication information you need to connect to that server. Click Forward to continue.
4. On the Configuration Complete window that appears, click Apply. At this point, the red icon turns blue and changes from an exclamation point to a check mark.
5. To register with Red Hat Network and check for software updates, proceed to the next section.

Registering with Red Hat Network

To register your system with Red Hat Network, you can click the blue RHN icon, use the `rhn_register` command or simply look for the Red Hat Network Configuration window to pop up when you run `up2date`. Using any of these methods requires a connection to the Internet and root access to your computer — the latter because the registration process needs to access your computer's RPM database. Though only the first computer you register can receive free automatic updates, you can register as many computers as you want with this tool. Here's how it works:

1. From a Terminal window type **`rhn_register`**. The first time you run this, the Red Hat Network Configuration window appears.
2. Select each of the following tabs to fill out basic configuration information:
   - **General** — Select an RHN server to use to get your updates. One should already be entered. If you need to go through an http proxy or provide a user name and password to reach the server, enter that information on this tab as well.
   - **Retrieval / Installation** — Set preferences here related to whether or not packages are immediately installed after they are downloaded and whether or not they are kept on your hard disk or deleted after installation. By default, packages are installed after retrieval, verified using GPG, and the RPMs are removed after they are installed. Packages that have had modifications to configuration files are not upgraded.
   - **Package Exceptions** — By default, any kernel packages are not upgraded. You can add other packages, or even selected files, to skip.
3. If you do not have the Red Hat, Inc. public key on your keyring (and you selected to use GPG as is set by default), you are prompted to install the key. Click Yes to install the key. The Welcome to Red Hat Update Agent window appears.
4. Click Forward. The Review the Red Hat Privacy Statement window appears.
5. Read the explanation of the registration procedure and click Next. The Red Hat Privacy Statement appears. This statement tells you, among other things, how your personal information is used and how Red Hat uses cookies to track activities.
6. Read the privacy statement. If it is acceptable, click Forward. The Login window appears.
7. Add information to create a user account and click Forward. The information you provide includes:
   - **User name** — Type a name of your choosing.
   - **Password** —Type a password (then type it again in the next box to verify it).
   - **E-mail address** — Type the address at which you would like Red Hat

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

to contact you.

The Register a User Account window appears.

Add the information and click Forward. The next screen asks for some personal information, including name, address, phone number, fax number, and how you would like to be contacted (e-mail, regular mail, phone, fax, or e-newsletter).

8. A Register a System Profile-Hardware window appears. This window displays information about your Red Hat Linux version, host name, IP address, CPU mode, CPU speed, and memory.
9. Select a profile name (the host name is used by default) or service ID number (which can be the computer's serial number) and click Forward.

The registration process begins building a list of packages installed on your computer. These are gathered so they can become part of your system profile. These packages are then displayed in the Register a System Profile-Packages window.

10. Check that the packages that appear in the window are all okay to include in your system profile (unselecting the ones you don't want to include) and click Next. The Send Profile Information to Red Hat Network window appears.
11. Click Forward to register the system-profile information with the Red Hat Network. (Or choose not to send the information by clicking Cancel.) Your information is sent and the Registration Finished window appears.
12. Click Forward to end the session.

After you have finished registering with Red Hat Network you can go to the RNH site and log in to use the service. You are now also ready to run the `up2date` command to update your software packages.

Getting Updates

After you have registered your computer with Red Hat Network you can use the `up2date` command (or right-click the notification icon and click Launch up2date) to find and install updates. As with `rhn_register`, you need root access to your computer to run `up2date`. This is because `up2date` needs permission to change configuration files and install packages.

The `up2date` configuration file (`/etc/sysconfig/rhn/up2date`) determines the behavior of `up2date`. Here are a few of the key settings in the `up2date` configuration file to help you understand what `up2date` will do when you launch it:

Debugging is disabled by default. Change `debug=0` to `debug=1` to turn on debugging if you are having trouble with `up2date`.

Downloaded packages are stored in the `/var/spool/up2date` directory. Once they are installed, the packages are deleted (by default).

The `up2date` log file is `/var/log/up2date`. You can open the file in any text editor to see descriptions of everything `up2date` did.

Any packages that begin with the word `kernel` are skipped. Add other package names you want skipped to the `pkgSkipList=kernel*` entry. You can use shell wildcard characters, such as the asterisk shown in the example, to match all kernel packages.

The `up2date` configuration file contains other settings that you can use to further refine your `up2date` session. For example, you can identify an HTTP proxy server, change the number of failed network-connection attempts that must occur before an exit, or indicate that you want to retrieve source packages along with the binary packages.

The up2date sources file (`/etc/sysconfig/rhn/sources`) is another configuration file you might consider making enhancements to. In particular, if you want to get updates of packages already installed on your system from yum or apt repositories, you can add the locations of those repositories to this file. The following two lines contain examples of entries you could add to the sources file for `yum` and `apt` repositories, respectively:

yum fedora-9-i386-stable http://download.fedora.us/fedora/redhat/9.0.93/i386/yum/stable

apt freshrpms-9-i386 http://ayho.freshrpms.net redhat/9/i386/ updates freshrpms

The `yum` example shows the service as `yum` and the channel name is `fedora-9-i386-stable`, followed by the location of the yum repository. In the `apt` example, the repository is divided into three parts: the host name, path, and final directory. Notice that in this example, both `updates` and `freshrpms` directories are available repositories. Once you have added a line similar to the ones just shown, they appear as channels in the Channel window when you run up2date.

To use the official locations for packages from Red Hat Inc., you can simply use the default `up2date`

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

settings. Here's an example of an `up2date` session:

1. Type **up2date** or click the notification icon and select Launch up2date. If your system does not yet have the Red Hat public key, you will be asked if you want to have `up2date` install it. (The key is needed to verify that the packages you are receiving are signed by Red Hat.)
2. Click Yes to install the key. The Red Hat Update Agent window appears.
3. Click Forward. The Channels window appears.
4. The Channel indicates the version of Red Hat Linux you are running. Select the Channel appropriate to your Red Hat installation and click Forward. The Update Agent begins checking your available packages.
5. The Available Package Updates window appears, displaying packages available for updates. Click each package you want to update or click Select All Packages to have all applicable packages updated. You can click on a package and then click View Advisory to see which advisories are addressed by updates for the package.
6. Click Forward to continue. The Update Agent begins testing packages to check for dependencies among the packages you are updating. If any dependencies arise from the packages you selected, the Packages Required to Solve Dependencies window appears, listing the additional packages you need.
7. Click Forward to have the packages shown included in the updates (or click the back button to change your package list). The Retrieving Packages window appears and notes the progress as the packages are downloaded from the Red Hat Network.
8. Click Forward to begin downloading the packages to your computer. (This may take a while, depending on the speed of your Internet connection and the number of packages you are downloading.)
9. When all packages are downloaded, click Forward. The Installing Packages window appears.
10. Click Forward to begin installing the packages you have downloaded.
11. When all the packages are installed, click Next. The All Finished window appears, displaying the package name and version for each package that was installed.
12. Click Finish to end the session.

If you are interested in seeing how the upgrades went, you can do the following:

Check the log file (`/var/log/up2date`).

See if any packages were left behind in the spool directory (`/var/spool/up2date`).

Though you don't necessarily have to reboot Linux at this point, if new daemon processes are installed with any of the updated packages you might want to restart them. For example, to restart your Web server (`httpd` daemon), you could type the following as root user:

/etc/init.d/httpd restart

Using the Red Hat Network

Once you have an account with Red Hat Network you can log in to the RHN site and take advantage of the services it offers. As noted earlier, each registered user is entitled to register one computer to receive automatic updates. Other services that you receive simply by registering with the Red Hat Network Web site include:

The ability to view complete listings of available errata, including security, bug fixes, and enhancement alerts. Errata are organized by Red Hat Linux version, to take the guesswork out of figuring out compatibility issues.

E-mail alerts of errata and notification of new packages available to fix the problem.

An entitlement manager that shows you which systems you are entitled to manage through RHN and lets you add more systems.

A feature for scheduling actions, such as automatic package update at a specified time.

Go to the Red Hat Network site (`rhn.redhat.com`) and log in using the user name and account you created during the `rhn_register` session described earlier. Traverse the RHN site using links on this page. Available sections include:

Your RHN (Contains notifications of the systems that may contain outdated packages. Sections in Your RNN include: Your Account (which contains personal information, such as name, address, e-mail, and so on) and Your Preferences (where you can specify whether or not you want to receive e-mail errata).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Systems (Lets you view a System List, which contains all the systems you have registered with RHN, or use a search tool to search your list of systems.
Errata (From the Errata page you can display all errata, show only those that apply to your system, or do keyword searches of available errata.
Software (A channel designates a version of Red Hat Linux that has package updates associated with it. By selecting the Channel List in this section, you can see what package versions are available for the Red Hat Linux version you are using.
Schedule (You can schedule package installations to occur at a later time. From this page you can also see pending, completed, and archived actions.
Help (For more help with Red Hat Network, the Help Desk offers the RHN FAQ, a customer-service support link, support forums, and the RHN Reference Guide.

If you want to purchase additional entitlements, I recommend that you check the Red Hat Network site for the latest prices. At the time of this writing, $60 per year Basic Service Entitlement is available for Fedora Core. This entitlement lets you get updated packages, but does not offer any technical support from Red Hat, Inc. (Previously, with Red Hat Linux, you could also get limited installation support, but that's not available with Fedora.) For official support, Red Hat is encouraging customers to upgrade to Red Hat Enterprise Linux and purchase the $96 per year Enterprise Service Entitlement.
Without purchasing additional entitlements, you can register additional systems with RHN. However, you can't use the `up2date` mechanism to update those additional systems. To register additional systems, simply run the `rhn_register` command from each system you want to register. When prompted, give RHN your existing user name and password. The next time you visit RHN you will see the newly registered system in your system list, although there will be no entitlement listed to let you do automatic updates.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

Although you may be using Red Hat Linux as a single-user system, many of the tasks you must perform to keep your computer running are defined as administrator tasks. A special user account called the root user is needed to do many of the things necessary to keep Linux working as you would like it to. If you are administering a Linux system that is used by lots of people, the task of administration becomes even larger. You must be able to add and support users, maintain the file systems, and ensure that system performance serves your users well.

To help the administrator, Linux comes with a variety of command-line utilities and graphical windows for configuring and maintaining your system. The `kudzu` program can be used to probe and reconfigure Red Hat Linux when you add or remove hardware. Commands such as `mkfs` and `mount` let you create and mount file systems, respectively. Tools like System Monitor and top let you monitor system performance.

Another important part of system administration is keeping up with the latest software fixes and enhancements to Red Hat Linux. Using the `up2date` facility, which is part of the Red Hat Network, you can automatically download and install updates to Red Hat Linux software packages on your system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 11: Setting Up and Supporting Users

Overview

One of the more fundamental tasks of administering a Red Hat Linux server is setting up and supporting user accounts. Computers, after all, are tools to be used by people. Apocalyptic science fiction plots aside, computers have no purpose without users.

When you install Red Hat Linux, you are required to create the root (administrator) user account. The first time you boot Red Hat Linux, you are asked to create a regular user account, using any name you choose. Several other administrative user accounts are set up automatically that you will probably never use directly.

Cross-Reference For a description of the root user account and how to use it, see Chapter 10.

This chapter discusses the basics of setting up a user account and offers tips on easing the burden of supporting a large number of Red Hat Linux users.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Creating User Accounts

Every person who uses your Red Hat Linux system should have a separate user account. Having a user account provides each person with an area in which to securely store files, as well as a means of tailoring his or her user interface (GUI, path, environment variables, and so on) to suit the way that he or she uses the computer.

You can add user accounts to your Red Hat Linux system in several ways. This chapter describes how to use the `useradd` command to add user accounts to Red Hat Linux from the command line, and how to use the Red Hat User Manager window to add users from the desktop.

Adding users with useradd

The most straightforward method for creating a new user from the shell is with the `useradd` command. After opening a Terminal window with root permission, you simply invoke the `useradd` command at the command prompt, with details of the new account as parameters.

The only required parameter to `useradd` is the login name of the user, but you will probably want to include some additional information. Each item of account information is preceded by a single letter option code with a dash in front of it. Table 11-1 lists the options that are available with the `useradd` command.

Table 11-1: useradd Command Options

Option	Description
`-c comment`	Provide a description of the new user account. Usually just the person's full name. Replace `comment` with the name of the user account.
`-d home_dir`	Set the home directory to use for the account. The default is to name it the same as the login name and to place it in `/home`. Replace `home_dir` with the directory name to use.
`-D`	Rather than create a new account, save the supplied information as the new default settings for any new accounts that are created.
`-e expire_date`	Assign the expiration date for the account in MM/DD/YYYY format. Replace `expire_date` with the expiration date to use.
`-f inactivity`	Set the number of days after a password expires until the account is permanently disabled. Setting this to `0` disables the account immediately after the password has expired. Setting it to `-1` disables the option, which is the default behavior. Replace `inactivity` with the number to use.
`-g group`	Set the primary group (as listed in the `/etc/group` file) that the new user will be in. Replace `group` with the group name to use.
`-G grouplist`	Add the new user to the supplied comma-separated list of groups.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 11-1: useradd Command Options

Option	Description
`-k skel_dir`	Set the skeleton directory containing initial configuration files and login scripts that should be copied to a new user's home directory. This parameter can only be used in conjunction with the `-m` option. Replace `skel_dir` with the directory name to use.
`-m`	Automatically create the user's home directory and copy the files in the skeleton directory (`/etc/skel`) to it.
`-M`	Do not create the new user's home directory, even if the default behavior is set to create it.
`-n`	Turn off the default behavior of creating a new group that matches the name and user ID of the new user.
`-o`	Use with `-u uid` to create a user account that has the same UID as another user name. (This effectively lets you have two different users with authority over the same set of files and directories.)
`-p passwd`	Enter a password for the account you are adding. This must be an encrypted password. Instead of adding an encrypted password here, you can simply use the `passwd user` command later to add a password for `user`.
`-r`	Allows you to create a new account with a user ID in the range reserved for system accounts.
`-s shell`	Specify the command shell to use for this account. Replace `shell` with the command shell.
`-u user_id`	Specify the user ID number for the account. The default behavior is to automatically assign the next available number. Replace `user_id` with the ID number.

As an example, let's create an account for a new user named Mary Smith with a login name of `mary`. First, log in as root, then type the following command:

useradd -c "Mary Smith" mary

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip When you choose a username, don't begin with a number (for example, 06jsmith). Also, it is best to use all lowercase letters, no control characters or spaces, and a ma

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Next, set Mary's initial password using the `passwd` command. It prompts you to type the password twice.

passwd mary

Changing password for user mary.

New password: *******

Retype new password: *******

Cross-Reference Refer to Chapter 14 for tips on picking good passwords.

In creating the account for Mary, the `useradd` command performs several actions:

- Reads the `/etc/login.defs` file to get default values to use when creating accounts.
- Checks command-line parameters to find out which default values to override.
- Creates a new user entry in the `/etc/passwd` and `/etc/shadow` files based on the default values and command-line parameters.
- Creates any new group entries in the `/etc/group` file.
- Creates a home directory based on the user's name and located in the `/home` directory.
- Copies any files located within the `/etc/skel` directory to the new home directory. This usually includes login and application startup scripts.

The preceding example uses only a few of the available `useradd` options. Most account settings are assigned using default values. Here is an example that uses a few more options:

useradd -m -g users -G wheel,sales -s /bin/tcsh -c "Mary Smith" mary

In this case, the `useradd` command is told to create a home directory for `mary` (`-m`), make `users` the primary group she belongs to (`-g`), add her to the groups `wheel` and `sales`, and assign `tcsh` as her primary command shell (`-s`). This results in a line similar to the following being added to the `/etc/passwd` file:

mary:x:502:100:Mary Smith:/home/mary:/bin/tcsh

In the `/etc/passwd` file, each line represents a single user account record. Each field is separated from the next by a colon (`:`) character. The field's position in the sequence determines what it is. As you can see, the login name is first. The password field contains an `x` because we are using a shadow password file to store encrypted password data. The user ID selected by the `useradd` command was 502. The primary group ID is 100, which corresponds to the `users` group in the `/etc/group` file. The comment field was correctly set to Mary Smith, the home directory was automatically assigned as `/home/mary`, and the command shell was assigned as `/bin/tcsh`, exactly as specified with the `useradd` options.

By leaving out many of the options (as I did in the first useradd example), defaults are assigned in most cases. For example, by not using `-g users` or `-G wheel,sales`, a group named mary would have been created and assigned to the new user. Likewise, excluding `-s /bin/tcsh`, causes `/bin/bash` to be assigned as the default shell.

The `/etc/group` file holds information about the different groups on your Red Hat Linux system and the users that belong to them. Groups are useful for allowing multiple people to share access to the same files while denying access to others. If you peek at the `/etc/group` file, you should find something similar to this:

bin:x:1:root,bin,daemon

daemon:x:2:root,bin,daemon

sys:x:3:root,bin,adm

adm:x:4:root,adm,daemon

tty:x:5:

disk:x:6:root

lp:x:7:daemon,lp

mem:x:8:

kmem:x:9:

wheel:x:10:root,joe,mary

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

.

.

.

nobody:x:99:

users:x:100:

chris:x:500

sheree:x:501

sales:x:601:bob,jane,joe,mary

Each line in the group file contains the name of a group, the group ID number associated with it, and a list of users in that group. By default, each user is added to his or her own group, beginning with GID 500. Note that `mary` was added to the `wheel` and `sales` groups instead of having her own group.

It is actually rather significant that `mary` was added to the `wheel` group. By doing this, you grant her the ability to use the `sudo` command to run commands as the root user (provided that `sudo` was configured as described in Chapter 10).

In this example, we used the `-g` option to assign `mary` to the `users` group. If you leave off the `-g` parameter, the default behavior is for `useradd` to create a new group with the same name and ID number as the user, which is assigned as the new user's primary group. For example, look at the following `useradd` command:

useradd -m -G wheel,sales -s /bin/tcsh -c "Mary Smith" mary

It would result in a `/etc/passwd` line like this:

mary:x:502:502:Mary Smith:/home/mary:/bin/tcsh

It would also result in a new group line like this:

mary:x:502:

Note that the user ID and group ID fields now have the same number. If you set up all of your users this way, you will have a unique group for every user on the system, which allows for increased flexibility in the sharing of files among your users.

Adding users with Red Hat User Manager

If you prefer a graphical window for adding, changing, and deleting user accounts, you can use the Red Hat User Manager window. To open the window from the GNOME desktop, click System Settings → Users and Groups (or type `redhat-config-users` from a Terminal window as root user). Figure 11-1 shows an example of that window.

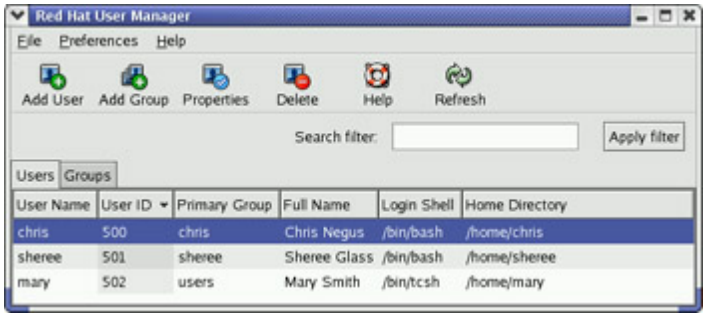

Figure 11-1: Manage users from the Red Hat User Manager window.

When you open the Red Hat User Manager window, you see a list of all regular users that are currently added to your computer. Administrative users (UID 1 through 100) are not displayed. For each user, you can see the user name, UID, primary group, full name, login shell, and home directory. Click on any of those headings to sort the users by that information.

To add a new user from the Red Hat User Manager window, do the following:

1. Click the Add User icon. A Create New User window appears, as shown in Figure 11-2.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Figure 11-2: Create New User window.

2. Type the requested information in the following fields:

   User Name (A single word to describe the user. Typically, the user name is eight characters, all lowercase, containing the user's real first name, last name, or (more often) a combination of the two (such as jwjones).

   Full Name (The user's full name (usually first name, middle initial, and last name). This name is typically just used for display, so using upper- and lowercase is fine.

   Password (The user's initial password. (Ask the user to change this password the first time he or she logs in to the new account, using the `passwd` command.)

   Confirm Password (Type the password again, to make sure you entered it correctly.

   Login Shell (The default shell (for entering typed commands) that the user sees when first logging in to Red Hat Linux from a character display.

   Create home directory (By default, this box is selected and the user's home directory (as indicated by the Home Directory field) is created automatically.

   Home Directory (By default, the user is given a home directory of the user's name in the `/home` directory. (For example, the user sheree would be assigned `/home/sheree` as her home directory.) Change this field if you want to assign the user to a different home directory.

   Create a private group for the user (Check this box if you want a group by the same name as the user, created for this user. The name is added to the `/etc/group` file. This feature is referred to as user private groups (UPGs).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip Using UPGs can be a benefit for sharing a directory of files among several users. Here's an example:

```
# useradd -m projectx
# mkdir /usr/local/x
# chown root.projectx /usr/local/x
# chmod 2775 /usr/local/x
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Specify user ID manually (Typically, you would not check this box, so that the UID for the new user would be assigned automatically. New UIDs for regular users start at 500. However, if you want to assign a particular UID for a user (for example, if you want to match the UID with the user's UID from another computer on your network), click this box and type the number you want to use in the UID box.

3. Click OK when you are done. The new user is added to the `/etc/passwd` and `/etc/group` files. The user account is now available for that user to login.

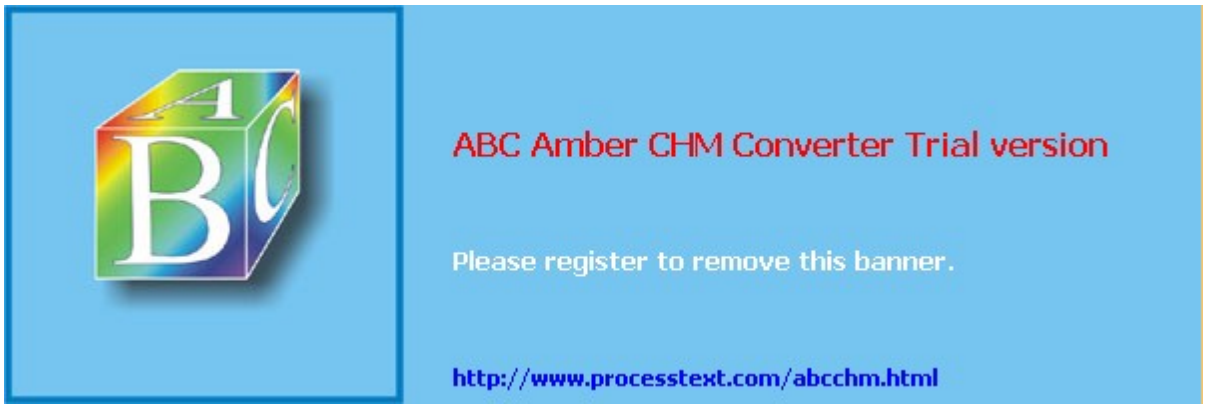

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Setting User Defaults

The `useradd` command and Red Hat User Manager window both determine the default values for new accounts by reading the `/etc/login.defs` file. You can modify those defaults by either editing that file manually with a standard text editor or by running the `useradd` command with the `-D` option. If you choose to edit the file manually, here is what you face:

```
# *REQUIRED*
# Directory where mailboxes reside, _or_ name of file, relative to the
# home directory. If you _do_ define both, MAIL_DIR takes precedence.
# QMAIL_DIR is for Qmail
#
#QMAIL_DIR Maildir
MAIL_DIR    /var/spool/mail
#MAIL_FILE .mail

# Password aging controls:
#
# PASS_MAX_DAYS Maximum number of days a password may be used.
# PASS_MIN_DAYS Minimum number of days allowed between password changes.
# PASS_MIN_LEN  Minimum acceptable password length.
# PASS_WARN_AGE Number of days warning given before a password
# expires.
#
PASS_MAX_DAYS    99999
PASS_MIN_DAYS    0
PASS_MIN_LEN     5
PASS_WARN_AGE    7

#
# Min/max values for automatic uid selection in useradd
#
UID_MIN                 500
UID_MAX               60000

#
# Min/max values for automatic gid selection in groupadd
#
GID_MIN                 500
GID_MAX               60000

#
# If defined, this command is run when removing a user.
# It should remove any at/cron/print jobs etc. owned by
# the user to be removed (passed as the first argument).
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
#
#USERDEL_CMD /usr/sbin/userdel_local

#
# If useradd should create home directories for users by default.
# On RH systems, we do. This option is ORed with the -m flag on
# useradd command line.
#
CREATE_HOME yes
```

Blank lines and comments beginning with a pound sign (#) are ignored. All other lines contain keyword/value pairs. For example, the keyword `MAIL_DIR` is followed by some white space and the value `/var/spool/mail`. This tells `useradd` that the initial user e-mail mailbox is created in that directory. Following that are lines that enable you to customize the valid range of automatically assigned user ID numbers or group ID numbers. A comment section that explains that keyword's purpose precedes each keyword. Altering a default value is as simple as editing the value associated with that keyword and then saving the `login.defs` file.

If you want to view the defaults, type the `useradd` command with the `-D` option as follows:

```
# useradd -D
GROUP=100
HOME=/home
INACTIVE=-1
EXPIRE=
SHELL=/bin/bash
SKEL=/etc/skel
```

You can also use the `-D` option to change defaults. When run with this flag, `useradd` refrains from actually creating a new user account; instead, it saves any additionally supplied options as the new default values in `/etc/login.defs`. Not all `useradd` options can be used in conjunction with the `-D` option. You may use only the five options listed in Table 11-2.

Table 11-2: useradd Options for Changing User Defaults

Options	Description
`-b default_home`	Set the default directory in which user home directories will be created. Replace `default_home` with the directory name to use. Usually this is `/home`.
`-e default_expire_date`	Set the default expiration date on which the user account is disabled. The `default_expire_date` value should be replaced with a date in the form MM/DD/YYYY — for example, 10/15/2001.
`-f default_inactive`	Set the number of days after a password has expired before the account is disabled. Replace `default_inactive` with a number representing the number of days.
`-g default_group`	Set the default group that new users will be placed in. Normally `useradd` creates a new group with the same name and ID number as the user. Replace `default_group` with the group name to use.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 11-2: useradd Options for Changing User Defaults

Options	Description
`-s default_shell`	Set the default shell for new users. Normally this is `/bin/sh`. Replace `default_shell` with the full path to the shell that you want as the default for new users.

To set any of the defaults, give the `-D` option first; then add any of the defaults you want to set. For example, to set the default home directory location to `/home/everyone` and the default shell to `/bin/tcsh`, type the following:

useradd -D -b /home/everyone -s /bin/tcsh

Besides setting up user defaults, an administrator can create default files that are copied to each user's home directory for use. These files can include login scripts and shell configuration files (such as `.bashrc`). The following sections describe some of these files.

Supplying initial login scripts

Many Red Hat Linux applications, including the command shell itself, read a configuration file at startup. It is traditional practice that these configuration files are stored in the users' home directories. In this way, each user can customize the behavior of the command shell and other applications without affecting that behavior for other users. In this way, global defaults can be assigned from `/etc/profile`, then those settings can be enhanced or overridden by a user's personal files.

The `bash` command shell, for example, looks for a file called `.bashrc` in the current user's home directory whenever it starts up. Similarly, the `tcsh` command shell looks for a file called `.tcshrc` in the user's home directory. You may see a repeating theme here. Startup scripts and configuration files for various applications usually begin with a dot (`.`) character and end in the letters `rc`. You can supply initial default versions of these and other configuration files by placing them in the `/etc/skel` directory. When you run the `useradd` command, these scripts and configuration files are copied to the new user's home directory.

Supplying an initial .bashrc file

By supplying your users with an initial `.bashrc` file, you give them a starting point from which they can further customize their shell environment. Moreover, you can be sure that the file is created with the appropriate access permissions so as not to compromise system security.

The `.bashrc` script is run each time the user starts a new `bash` shell. So, security is a concern. It is also a good place to supply useful command aliases and additions to the command search path. Here's an example:

```
# .bashrc

# User specific aliases and functions

alias rm='rm -i'

alias cp='cp -i'

alias mv='mv -i'


if [ -f /etc/bashrc ]; then

        . /etc/bashrc

fi


PATH=$PATH:/usr/bin:/usr/local/bin

export PATH
```

This sample `.bashrc` file creates aliases for the `rm`, `cp`, and `mv` commands that result in a `-i` option always being used (unless overridden with the `-f` option). This protects against the accidental deletion of files. Next, the file executes the `/etc/bashrc` (if it exists) to read any further global bash values. This file also sets the search path.

Supplying an initial .tcshrc file

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

This following example `.tcshrc` file does basically the same thing as the preceding `.bashrc` example. However, this file (which is for the root user) has the additional task of setting the appearance of the command prompt:

```
# .tcshrc

# User specific aliases and functions

alias rm 'rm -i'
alias cp 'cp -i'
alias mv 'mv -i'

setenv PATH "$PATH:/usr/bin:/usr/local/bin"

set prompt='[%n@%m %c]# '
```

Instead of using the `export` command to set environment variables, the `tcsh` shell uses the `setenv` command. In the example, `setenv` is used to set the `PATH` variable. The shell prompt is set to include your user name (`%n`), your computer name (`%m`), and the name of the current directory (`%c`). So, if you were to use the `tcsh` shell as the root user on a computer named `maple` with `/tmp` as your current directory, your prompt would appear as follows:

```
[root@maple /tmp]#
```

The `.tcshrc` file can also be named `.cshrc`. The `tcsh` shell is really an extended version of the `csh` shell (in fact, you can invoke it by the `csh` name). When a `tcsh` shell is started, it first looks for a `.tcshrc` file in the current user's home directory. If it can't find a file by that name, it looks for the other name, `.cshrc`. Thus, either name is appropriate.

Configuring system-wide shell options

Allowing individually customizable shell startup files for each user is a very flexible and useful practice. But sometimes you need more centralized control than that. You may have an environment variable or other shell setting that you want set for every user, without exception. If you add that setting to each individual shell, the user has the ability to edit that file and remove it. Furthermore, if that setting must be changed in the future, you must change it in every single user's shell startup file.

Fortunately, there is a better way. There are default startup files that apply to all users of the computer that each command shell reads before reading the user-specific files. In the case of the `bash` command shell, it reads the `/etc/bashrc` file before doing anything else.

Similarly, the `tcsh` shell reads the `/etc/csh.cshrc` file before processing the `.cshrc` or `.tcshrc` file found in the user's home directory. The following `/etc/csh.cshrc` file ships with Red Hat Linux:

```
# /etc/cshrc
#
# csh configuration for all shell invocations.

# by default, we want this to get set.
# Even for non-interactive, non-login shells.
 [ `id -gn` = `id -un` -a `id -u` -gt 99 ]
if $status then
    umask 022
else
    umask 002
endif
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
if ($?prompt) then
  if ($?tcsh) then
    set prompt='[%n@%m %c]$ '
  else
    set prompt=\[`id -nu`@`hostname -s`\]\$\
  endif
endif
```

The `/etc/cshrc` and `/etc/bashrc` files set a variety of shell environment options. If you wish to modify or add to the shell environment supplied to every single user on the system, the `/etc/bashrc` or `/etc/cshrc` files are the place to do it.

Setting system profiles

Some of the most basic information assigned to each user is added from the `/etc/profile` file. So, if you want to change any of the following information, you can start from `/etc/profile`. Here are some values contained in `/etc/profile`:

PATH (Assigns the default PATH for the root user and for all other users. You might change this value to add paths to local directories containing applications all users need.

ulimit (Sets the maximum allowable file size the user can create from the shell to be unlimited. You can use ulimit to restrict maximum file size if you find that users are creating enormous files. As defined in the `/etc/profile` file, `ulimit` sets no limit to the size of files a user can create. However, it does prevent core files (normally created when a process crashes) from being created.

Environment variables (Shell environment variables that are needed for standard operation are assigned in this file. These include USER (set by the `id –un` command), LOGNAME (same as USER), MAIL (set to `/var/spool/mail/$USER`), HOSTNAME (set to `/bin/hostname`), and HISTSIZE (which sets shell command history to 1000 items).

INPUTRC (Sets keyboard mappings for particular situations, based on the contents of the `/etc/inputrc` file. In particular, the inputrc file makes sure that the Linux console and various Terminal windows (`xterm` and `rxvt`) all behave sanely.

The last thing that the `/etc/profile` file does is look at the contents of the `/etc/profile.d` directory and source in the files that it finds. Each file contains settings that define environment variables or aliases that affect how users can use the shell. For example, the `lang.sh` and `lang.csh` files identify the locations of foreign language files. The `vim` files create aliases that cause `vim` to be used when `vi` is typed. The `which-2.sh` file defines a set of options used by the `which` command. You can modify the `profile.d` files or add your own to have environment variables and aliases set for all of your users.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Creating Portable Desktops

Linux is an operating system that was born on the Internet, so it is not surprising that it has strong networking capabilities. This makes Linux an excellent server, but it also allows Linux to be an excellent desktop workstation, especially in a highly networked environment. Red Hat Linux lets you easily set up your users with a portable desktop that follows them from computer to computer. With other leading desktop operating systems, it is not nearly as easy.

Normally, a Red Hat Linux user's home directory is located within the `/home` directory. As an alternative, within the home directory you can create a directory named after the system's host name. Within that directory, create the users' home directories. Thus, on a Linux system named `dexter`, the user `mary` would have a home directory of `/home/dexter/mary` instead of `/home/mary`. There is a very good reason for doing this.

If you are logged into the Linux system `ratbert` and wish to access your home directory on `dexter` as if it were stored locally, the best approach is to use Network File System (NFS) to mount `dexter`'s `/home` directory on the `/home` on `ratbert`. This results in having the same contents of your home directory available to you no matter which machine you log in to.

Cross-Reference You can read more about NFS in Chapter 18.

To mount `dexter's /home` directory as described, you would add a line similar to the following in `ratbert's /etc/fstab` file:

```
dexter:/home /home nfs defaults 0 0
```

You would also add an entry such as the following in `dexter's /etc/exports` directory:

```
/home ratbert
```

Now, when `ratbert` boots up, it automatically mounts `dexter's` home partition over the network. This enables us to treat the remote files and directories on `dexter's /home` as if they are locally stored on `ratbert`. Unfortunately, this has the side effect of "covering up" `ratbert's` actual `/home` directory.

This is where the extra directory level based on the system name comes to the rescue. With all of `dexter's` home directories located in `/home/dexter` and all of `ratbert's` home directories located in `/home/ratbert`, we can remove the danger of one system covering up the home directories of another. In fact, let's take this example one step further: Imagine a scenario in which the systems `dexter`, `ratbert`, and `daffy` all have portable desktops that are shared with the other systems. The `/etc/fstab` and `/etc/exports` files for each system should have the following lines added to them.

The `/etc/exports` and `/etc/fstab` files for `dexter` are as follows:

/etc/exports file

```
/home/dexter ratbert,daffy
```

/etc/fstab file

```
Ratbert:/home/ratbert /home/ratbert nfs defaults 0 0

Daffy:/home/daffy    /home/daffy   nfs defaults 0 0
```

The `/etc/exports` and `/etc/fstab` files for `ratbert` are:

/etc/exports

```
/home/ratbert dexter,daffy
```

/etc/fstab

```
dexter:/home/dexter /home/dexter nfs defaults 0 0

daffy:/home/daffy   /home/daffy  nfs defaults 0 0
```

The `/etc/exports` and `/etc/fstab` files for `daffy` are:

/etc/exports

```
/home/dexter ratbert,dexter
```

/etc/fstab

```
Ratbert:/home/ratbert /home/ratbert  nfs defaults 0 0

Dexter:/home/dexter   /home/dexter   nfs defaults 0 0
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

As you can see, each system uses NFS to mount the home directories from the other two systems. A user can travel from server to server and see exactly the same desktop on each system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Providing Support to Users

Creating new user accounts is just one small administrative task among many. No single chapter can adequately discuss all the tasks that are involved in the ongoing support of users. But I share with you a few hints and procedures to ease that burden.

Creating a technical support mailbox

E-mail is a wonderful communication tool, especially for the overworked system administrator. People usually put more thought and effort into their e-mail messages than into the voice messages that they leave. A text message can be edited for clarity before being sent, and important details can be cut and pasted from other sources. This makes e-mail an excellent method for Red Hat Linux users to communicate with their system administrator.

In an office with only a few users, you can probably get away with using your personal mailbox to send and receive support e-mails. In a larger office, however, you should create a separate mailbox reserved only for technical support issues. This has several advantages over the use of your personal mailbox:

- Support messages will not be confused with personal, non-support-related messages.
- Multiple people can check the mailbox and share administrative responsibility without needing to read each other's personal e-mail.
- Support e-mail is easily redirected to another person's mailbox when you go on vacation. Your personal e-mail continues to go to your personal mailbox.

One easy solution is to simply create a support e-mail alias that redirects messages to an actual mailbox or list of mailboxes. For example, suppose you wish to create a support alias that redistributes e-mail to the user accounts for support staff members Joe, Mary, and Bob. You would log in as root, edit the `/etc/aliases` file, and add lines similar to the following:

```
# Technical support mailing list

support: joe, mary, bob
```

After saving the file, you need to run the `newaliases` command to recompile the `/etc/aliases` file into a database format. Now your users can send e-mail to the support e-mail address, and the message is automatically routed to everyone on the list. When a member of the list responds to that message, he or she should use the "Reply To All" option so that the other support staff members also see the message. Otherwise, multiple people may attempt to solve the same problem, resulting in wasteful duplication of effort.

You may also choose to create a support user account. The technical support staff would log in to this account to check messages and send replies. In this manner, all replies are stamped with the support login name and not the personal e-mail address of a staff member.

Resetting a user's password

One common (if not *the* most common) problem that your users will encounter is the inability to log in because:

- They have the Caps Lock key on.
- They have forgotten the password.
- The password has expired.

If the Caps Lock key is not on, then you probably need to reset the individual's password. You can't look up the password because Red Hat Linux stores passwords in an encrypted format. Instead, use the `passwd` command to assign a new password to the user's account. Tell the user what that new password is (preferably in person), but then set the password to expire soon so that he or she must choose one (hopefully, a new one that is more easily remembered).

If you must reset a user's password, do so with the `passwd` command. While logged in as root, type `passwd` followed by the login name you are resetting. You are prompted to enter the password twice. For example, to change the password for `mary`, type:

```
# passwd mary
```

After resetting the password, set it to expire so that the user is forced to change it the next time she logs in. You can use the `chage` command to set an expiration period for the password and to trick the system into thinking that the password is long overdue to be changed.

```
# chage -M 30 -d 0 mary
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The `-M 30` option tells the system to expire Mary's password every 30 days. The `-d 0` option tricks the system into thinking that her password has not been changed since January 1, 1970.

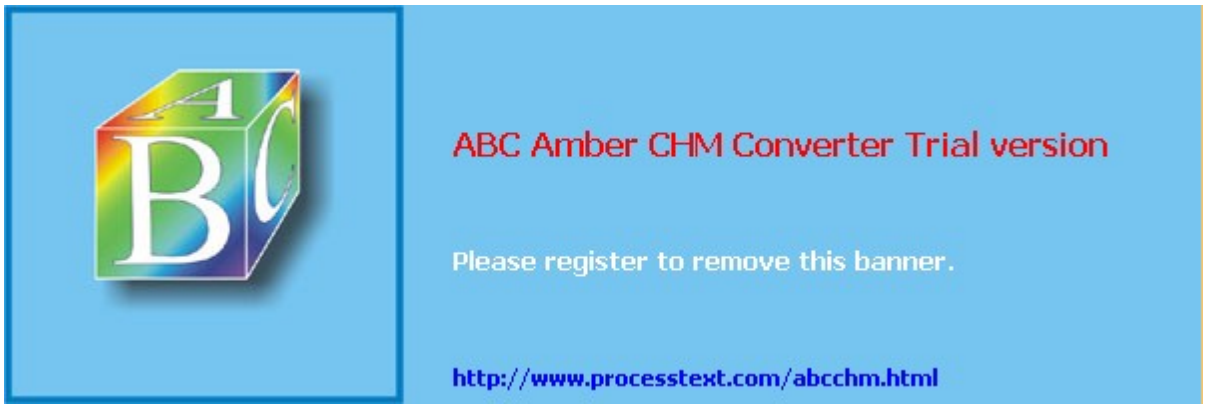

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Modifying Accounts

Occasionally, a user needs more done to an account than just a resetting of the password. You may need to change the groups that user is in, or the drive that a home directory resides on. The following sections explain how to modify user accounts using one of two methods: `usermod` or the Red Hat User Manager.

Modifying user accounts with usermod

The `usermod` command is similar to the `useradd` command and even shares some of the same options. However, instead of adding new accounts, it enables you to change various details of existing accounts. When invoking the `usermod` command, you must provide account details to change followed by the login name of the account. Table 11-3 lists the available options for the `usermod` command.

Table 11-3: usermod Options for Changing Existing Accounts

Options	Description
`-c comment`	Change the description field of the account. You can also use the `chfn` command for this. Replace `comment` with a name or other description of the user account, placing multiple words in quotes.
`-d home_dir`	Change the home directory of the account to the specified new location. If the `-m` option is included, copy the contents of the home directory as well. Replace `home_dir` with the full path to the new directory.
`-e expire_date`	Assign a new expiration date for the account, replacing `expire_date` with a date in MM/DD/YYYY format.
`-f inactivity`	Set the number of days after a password expires until the account is permanently disabled. Setting `inactivity` to `0` disables the account immediately after the password has expired. Setting it to `-1` disables the option, which is the default behavior.
`-g group`	Change the primary group (as listed in the `/etc/group` file) that the user is in. Replace `group` with the name of the new group.
`-G grouplist`	Set the list of groups that user belongs to. Replace `grouplist` with a list of groups.
`-l login_name`	Change the login name of the account to the name supplied after the `-l` option. Replace `login_name` with the new name. This automatically changes the name of the home directory; use the `-d` and `-m` options for that.
`-m`	This option is used only in conjunction with the `-d` option. It causes the contents of the user's home directory to be copied to the new directory.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 11-3: usermod Options for Changing Existing Accounts

Options	Description
`-o`	This option is used only in conjunction with the `-u` option. It removes the restriction that user IDs must be unique.
`-s shell`	Specify a new command shell to use with this account. Replace `shell` with the full path to the new shell.
`-u user_id`	Change the user ID number for the account. Replace `user_id` with the new user ID number. Unless the `-o` option is used, the ID number must not be in use by another account.

Assume that a new employee named Jenny Barnes will be taking over Mary's job. We want to convert the `mary` account to a new name (`-l jenny`), new comment (`-c "Jenny Barnes"`), and home directory (`-d /home/jenny`). We could do that with the following command:

usermod -l jenny -c "Jenny Barnes" -m -d /home/jenny mary

Furthermore, if after converting the account we learn that Jenny prefers the `tcsh` shell, we could make that change with the `-s` option (`-s /bin/tcsh`):

usermod -s /bin/tcsh jenny

Instead, we could use the `chsh` command to change the shell. The following is an example:

chsh -s /bin/tcsh jenny

The `chsh` command is handy because it enables a user to change his or her own shell setting. Simply leave the user name parameter off when invoking the command, and `chsh` assumes the currently logged-in user as the account to change.

Modifying user accounts with Red Hat User Manager

To use the desktop to change an existing account, you can use the Red Hat User Manager. Here's how to add a new user from the Red Hat User Manager window:

From the Red Hat menu, click on System Settings→ Users and Groups (or type `redhat-config-users` from a Terminal window as root user). The main Red Hat User Manager window appears.

Click on the user name of the account you want to modify, then select the Properties button. A User Properties window appears, as shown in Figure 11-3.

Figure 11-3: Choose Properties to modify an existing user account.

There are four tabs of information you can modify for the user you selected:

User Data (This tab contains the user information you created when you first added the user account.

Account Info (Click on "Enable account expiration" and type a date if you want the account to become inaccessible after a particular date. Click on "User account is locked" if you want to prevent access to the account but not delete it. (The latter is a good technique when an employee is leaving the company or if you want to lock out a customer whose account is temporarily disabled. The information isn't removed, it

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

just isn't accessible.)

Password Info (Click on "Enable password expiration" if you want to control expiration of the user's password. By default, passwords don't expire. Here are your options: "Days before change allowed" (forces the user to keep the password for at least a set number of days before it can be changed); "Days before change required" (allows the user to keep the same password for at least the set number of days); "Days warning before change" (sets how many days before the password expiration day that the user is warned to change the password); "Days before account inactive" (sets the number of days after which the account is deactivated).

Groups (Select from the list of available groups to add the user to one or more of those groups.

Click Apply to apply the changes to the user account.

Team LiB

PREVIOUS NEXT

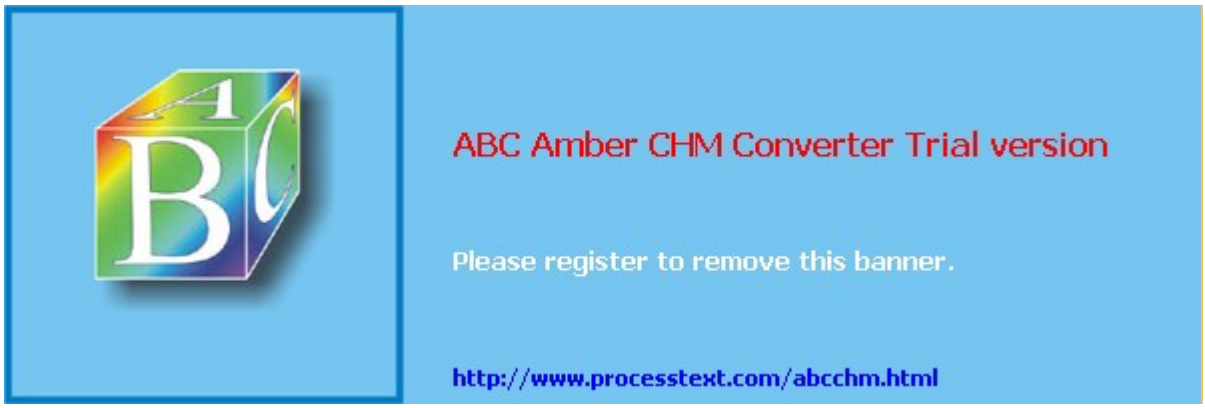

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Deleting User Accounts

Occasionally, it is necessary to remove a user account from your Red Hat Linux system. This can be done with either the `userdel` command or the Red Hat User Manager window.

Deleting user accounts with userdel

The `userdel` command takes a single argument, which is the login name of the account to delete. If you supply the optional `-r` option, it also deletes the user's home directory and all the files in it. To delete the user account with login name `mary`, you would type this:

userdel mary

To wipe out her home directory along with her account, type this:

userdel -r mary

Files owned by the deleted user but not located in the user's home directory will not be deleted. The system administrator must search for and delete those files manually. The `find` command comes in very handy for this type of thing. I won't describe all the capabilities of the `find` command (that would take a very fat chapter of its own). I do, however, provide a few simple examples of how to use `find` to locate files belonging to a particular user, even when those files are scattered throughout a file system. You can even use the `find` command to delete or change the ownership of files, as they are located. Table 11-4 has a few examples of the `find` command in action.

Table 11-4: Using find to Locate and Change User Files

Find Command	Description
`find / -user mary`	Search the entire file hierarchy (start at `/`) for all files and directories owned by `mary` and print the filenames to the screen.
`find /home -user mary -exec rm -i {} \;`	Search for all files and subdirectories under `/home` owned by `mary`. Run the `rm` command interactively to delete each file.
`find / -user mary -exec chown jenny {} \;`	Search for all files and subdirectories under `/home` that are owned by user `mary` and run the `chown` command to change each file so that it is owned by `jenny` instead.
`find / -uid 500 -exec chown jenny {} \;`	This command is basically the same as the previous example, but it uses the user ID number instead of the user name to identify the matching files. This is useful if you have deleted a user before converting her files.

There are a few common things about each invocation of the `find` command. The first parameter is always the directory to start the recursive search in. After that come the file attributes to match. You can use the `-exec` parameter to run a command against each matching file or directory. The `{}` characters designate where the matching filename should be filled in when `find` runs the `-exec` option. The `\;` at the end simply tells Linux where the command ends. These are only a few of `find's` capabilities. I encourage you to read the online man page to learn more about `find`. (Type **man find** to view the page.)

Deleting user accounts with Red Hat User Manager

To delete a user from the Red Hat User Manager window, simply click the line representing the user account, then click the Delete button.

The information about the user is removed from the `/etc/passwd` file; thus, the user can no longer log in.

The home directory and all files owned by the user will still exist. However, a listing of files previously owned by that user (`ls -l`) will show only the former user's UID, but

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

no name, as the owner.

See the description in the previous section for information about how to find and remove files previously owned by the user.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip You may want to transfer the ownership of the files from the old user to the new user (if, for example, a new employee is taking over the work of an employee that is leaving the co

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Checking Disk Quotas

Limited disk space can be another source of user support calls. Red Hat Linux offers the quotas software package for limiting and displaying the amount of disk space that a user can consume. You can also use the `du` command to see how much disk space has been used in a particular directory (and related subdirectories). To automate the process of checking for disk space, you can create your own script. The following sections describe these ways of dealing with potential disk space problems.

Using quota to check disk usage

A careless or greedy user can gobble up all the space on your hard disk and, possibly, bring your computer to a halt. By using disk quotas, you can limit the amount of disk resources a user or group can use up.

The quota package contains a set of tools that lets you limit the amount of disk space (based on disk blocks) and files (based on inodes) that a user can consume. Using quotas, you can limit the amount of usage (on a per-user and -group basis) for each file system on your computer. The general steps for setting disk quotas are:

- Creating quota files
- Editing the `/etc/fstab` file
- Creating quota rules
- Checking quotas

You set quotas on disk partitions listed in your `/etc/fstab` file. For computers that are shared by many users, there may be a separate `/home` or `/var` partition where users are expected to put all their data. That kind of partition boundary can prevent an entire disk from being consumed by one user. Quotas on the `/home` or `/var` partition can make sure that the space within those partitions are shared fairly among your computer's users.

The procedure that spans the next few sections assumes that you have a separate `/home` partition on your computer for which you want to create quotas. You could use any partition, not just the `/home` partition shown in the procedure. For example, if you have only one partition mounted at the root of the files system (`/`), you could set quotas for your entire file system by replacing `/home` with `/` in the following example.

Editing the /etc/fstab file

You need to add quota support to the file system. To do that, edit the `/etc/fstab` file and add the `usrquota` option to field number four of the partition for which you want to set quotas. Here is an example of a line from `/etc/fstab`:

```
/dev/hda2    /home    ext3    defaults,usrquota,grpquota    1 2
```

Here, the `/home` file system is used to allow disk quotas for all users' home directories under the `/home` directory.

Before the `usrquota` option can take effect, the file system must be remounted. This happens automatically when you reboot, which you will have to do if you are setting quotas for the root (/) file system. Otherwise, you might be able to use the `umount` and `mount` commands to cause the `usrquota` option to take effect.

Creating quota files

You need to have `aquota.user` and/or `aquota.group` files in the root directory of the partition on which you want to establish disk quotas. To add quotas based on individual users, you need an `aquota.user` file, while `aquota.group` is needed to set quotas based on groups. One way to create these files is with the `quotacheck` command. Here is an example of the `quotacheck` command to create an initial `aquota.user` file:

```
# quotacheck -c /home
```

A `/home/aquota.user` file is created from the previous command. (To create an initial `aquota.group` file, type `touch /home/aquota.group`.) Next, you must create the disk usage table for the partition. Here's an example of how to do that:

```
# quotacheck -vug /home
```

The `quotacheck` command in this example looks at the file system partition mounted on `/home` and

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

builds a table of disk usage. The `-v` option produces verbose output from the command, the `-u` option causes user quotas to be examined, and the `-g` option causes group quotas to be examined. Permissions on the two files are set so only root can access them (`chmod 600 aquota.*`).

Creating a quota startup script

If the quota package doesn't include a startup script (and it doesn't with the current Red Hat Linux), you can create your own. You want this script to check quotas (`quotacheck` command), start the quota service (`quotastart` command) and turn off the service (`quotaoff` command).
Open a new file called `/etc/init.d/quota` as root user, using any text editor. Here is an example of the content you can add to that file:

```
#!/bin/bash


# init file for quota
#
# description: Checks disk quotas
#
# processname: quota
# chkconfig: - 90 90
# source function library
. /etc/rc.d/init.d/functions


case "$1" in
  start)
     echo -n "Checking quotas: "
        daemon /sbin/quotacheck -avug
     echo
     echo -n "Starting quotas: "
        daemon /sbin/quotaon -avug
     echo
     ;;
  stop)
     echo -n "Shutting down quotas: "
     daemon /sbin/quotaoff -a
     echo
     ;;
  restart)
        $0 stop
        $0 start
        ;;
  *)
     echo "Usage: quota {start|stop|restart}"
     exit 1
esac


exit 0
```

The quota script, when started, first runs the `quotacheck` command to check all file systems for which quota checking is on. Then it turns on quota checking with the `quotaon` command.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Turn on the quota startup script

If you created a quota file, as described in the previous step, you need to make it executable and set it to start automatically when you start Red Hat Linux. To do those things, type the following as root user:

chmod 755 /etc/init.d/quota

chkconfig on quota

At this point, links are created so that your quota script starts when Red Hat Linux boots.

Creating quota rules

You can use the `edquota` command to create quota rules for a particular user or group. (Valid users and groups are listed in the `/etc/passwd` and `/etc/group` files, respectively). Here is an example of an `edquota` command to set quotas for a user named jake.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The edquota command uses the vi text editor to edit your quota files. To use a different editor, change the value of the EDITOR or VISUAL environment variable before running edquota. For example, to use the emacs editor, type the following before running edquota:

```
# export EDITOR
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

edquota -u jake

Disk quotas for user jake (uid 501)

Filesystem	blocks	soft	hard	inodes	soft	hard
/dev/hda2	596	0	0	1	0	0

~

~

~

"/tmp//EdP.aBY1zYC" 3L, 215C

This example shows that user quotas can be set for the user jake on the `/dev/hda2` partition (which is `/home` in our example). Currently, jake has used 596 blocks (a block equals 1K on this `ext3` file system). One file was created by jake (represented by 1 inode). To change the disk usage limits, you can edit the zeros (unlimited use) under the soft and hard heading for blocks and inodes.

Soft limits set limits that you don't want a user or group to exceed. Hard limits set the boundaries that you will not let a user or group exceed. After a set grace period that a soft limit is exceeded (which is seven days, by default), the soft limit becomes a hard limit. (Type **`edquota -t`** to check and change the grace periods that you have set.)

Here is an example of how the line in the previous `edquota` example could be changed:

/dev/hda2 596 512000 716800 1 800 1000

In this example, the soft limit on the number of blocks that the user jake could consume on the `/dev/hda2` device (`/home`) is `512000` blocks (or 500MB) the hard limit is `716800` blocks (or 700MB). Soft and hard limits on inodes are 800 and 1000, respectively. If either of the soft limits are exceed by the user jake, he has seven days to get back under the limit, or he will be blocked from using any more disk space or inodes.

Further attempts to write to a partition after the hard limit has been exceeded results in a failure to write to the disk. When this happens, the user that tries to create the file that exceeds his limit will see the a message like the following:

ide0(3,2): write failed, user block limit reached.

cp: writing 'abc.doc': Disk quota exceeded

Instead of assigning quotas to users, you can assign quotas to any group listed in the `/etc/group` file. Instead of the `-u` option to `edquota`, use the `-g` options followed by a group name.

Updating quota settings

After you have changed quota settings for a user, you should rerun the `quotacheck` command. You should also run the `quotacheck` command periodically, to keep the quota records up to date. One way to do that is to run the `quotacheck` command weekly using a cron entry.

Checking quotas

To report on how much disk space and how many inodes each user on your computer (for which you have set quotas) has consumed, use the `repquota` command. Here is an example of the `repquota` command for reading quota data relating to all partitions that are using quotas:

repquota -a

*** Report for user quotas on device /dev/hda2

Block grace time: 7days: Inode grace time: 7days

	Block limits				File limits			
User	used	soft	hard	grace	used	soft	hard	grace
root --	1973984	0	0		2506	0	0	
jake --	1296	700	1700	6days	3	0	0	

In this previous example, jake has exceeded his soft limit of 700 blocks. He currently has six days left in his grace period to remove enough files so that the soft limit does not become the hard limit.

Using du to check disk use

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You can discover the most voracious consumers of disk space using the `du` command. Invoke `du` with the `-s` option and give it a list of directories; it reports the total disk space used by all the files in each directory. Add an -h option to display disk space used in numbers, followed by kilobytes (k), megabytes (M), or gigabytes (G). The -c option adds a total of all requested directories at the end. The following checks disk usage for several home directories:

du -h -c -s /home/tom /home/bill /home/tina /home/sally

This should result in a list of all of your users' home directories preceded by the number of kilobytes that each directory structure uses. It looks something like this:

339M /home/tom

81M /home/bill

31M /home/tina

44k /home/sally

450M total

Removing temp files automatically

Some potential disk-consumption problems are set up to take care of themselves. For example, directories for storing temporary files used by applications (such as `/tmp` and `/var/tmp`) can consume lots of disk space over time. To deal with the problem, Red Hat Linux includes the tmpwatch facility. The `/usr/sbin/tmpwatch` command runs from the cron file `/etc/cron.daily/tmpwatch` to delete unused temp files. Here's what that file contains:

```
/usr/sbin/tmpwatch 240 /tmp
/usr/sbin/tmpwatch 720 /var/tmp
for d in /var/{cache/man,catman}/{cat?,X11R6/cat?,local/cat?}; do
    if [ -d "$d" ]; then
     /usr/sbin/tmpwatch -f 720 $d
    fi
done
```

Each day, this tmpwatch scripts runs to delete temporary files that haven't been used for a while. Files from the `/tmp` and `/var/tmp` directories are removed after 240 and 720 hours of not being accessed, respectively. Temporary man page files stored in `/var/cache` subdirectories are also checked and deleted after 720 hours of disuse.

Team LiB

PREVIOUS NEXT

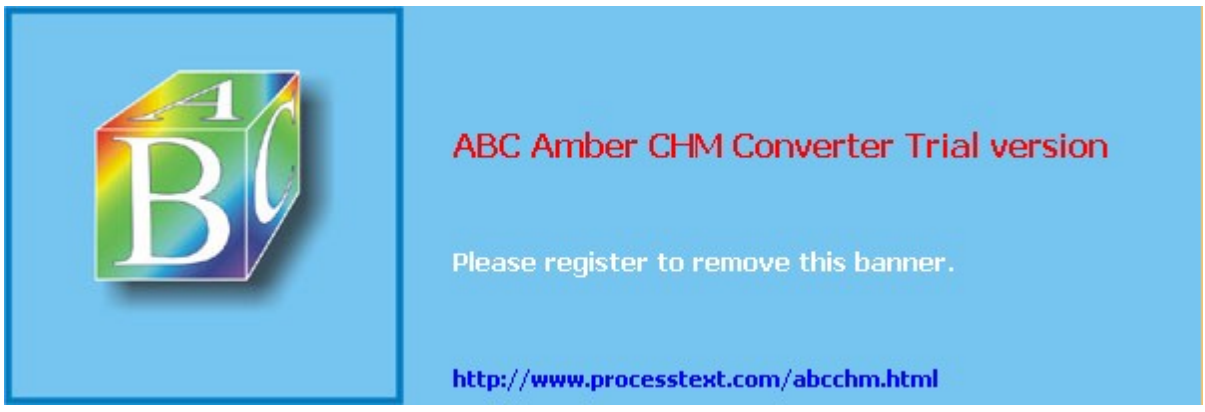

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Sending Mail to All Users

Occasionally, you need to send messages to all users on your system. Warning users of planned downtime for hardware upgrades is a good example. Sending e-mail to each user individually is extremely time consuming and wasteful; this is precisely the kind of task that e-mail aliases and mailing lists were invented for. Keeping a mailing list of all the users on your system can be problematic, however. If you are not diligent about keeping the mailing list current, it becomes increasingly inaccurate as you add and delete users. Also, if your system has many users, the mere size of the alias list can become unwieldy.

The following script, called `mailfile`, provides a simple method of working around these problems. It grabs the login names from the `/etc/passwd` file and sends e-mail to all users.

```
#!/bin/csh
#
# mailfile: This script mails the specified file to all users
#           of the system.  It skips the first 17 accounts so
#           we do not send the email to system accounts like
#           'root'.
#
# USAGE: mailfile "Subject goes here" filename.txt


#
# Check for a subject
#
if ( `echo $1 | awk '{ print $1 }'` == "" ) then
   echo You did not supply a subject for the message.
   echo Be sure to enclose it in quotes.
   exit 1
else
   # Get the subject of the message
   set subject=$1
endif


#
# Check for a filename
#
if ( $2 == "" ) then
   echo You did not supply a file name.
   exit 2
else
   # Get the name of the file to send
   set filename=$2
endif


#
# Check that the file exists
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
#
if ( -f $filename ) then
  echo Sending file $filename
else
  echo File does not exist.
  exit 3
endif


#
# Loop through every login name, but skip the first 17 accounts
#
foreach user ( `awk -F: '{ print $1 }' /etc/passwd | tail +17` )
   # Mail the file
   echo Mailing to $user
   mail -s "$subject" $user < $filename


   # sleep for a few seconds so we don't overload the mailer
   # On fast systems or systems with few accounts, you can
   # probably take this delay out.
   sleep 2
end
```

The script accepts two parameters. The first is the subject of the e-mail message, which is enclosed in quotes. The second is the name of the file containing the text message to send. Thus, to send an e-mail message to all users warning them about an upcoming server hardware upgrade, I may do something similar to the following:

```
mailfile "System upgrade at 5:00pm" upgrade.txt
```

The file `upgrade.txt` contains the text of the message to be sent to each user. The really useful thing about this approach is that I can save this text file and easily modify and resend it the next time I upgrade the system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If your users log in to your system using text-based logins instead of graphical logins, you can add messages to the /etc/motd file to have them reach your users. Any text in that

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

It is not uncommon for a Red Hat Linux system to be used as a single-task server with no actual users. It sits quietly in a server room, serving Web pages or handling domain name service, never crashing, and rarely needing attention. This is not always the case, however. You may have to support users on your Red Hat Linux server, and that can be the most challenging part of your system administration duties.

Red Hat Linux provides a variety of tools that help you with your administrative chores. The `useradd`, `usermod`, and `userdel` commands enable easy command-line manipulation of user account data. Furthermore, creating a support mailbox and building shell scripts to automate repetitive tasks lightens your load even more. Red Hat Linux builds on top of the rich history of UNIX and provides an ideal platform to support the diverse needs of your users.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 12: Automating System Tasks

Overview

You'd never get any work done if you typed every command that needs to be run on your Red Hat Linux system when it starts. Likewise, you could work more efficiently if you grouped together sets of commands that you run all the time. Shell scripts can handle these tasks.
A *shell script* is a group of commands, functions, variables, or just about anything else you can use from a shell. These items are typed into a plain-text file. That file can then be run as a command. Red Hat Linux uses system initialization shell scripts during system start-up to run commands needed to get things going. You can create your own shell scripts to automate the tasks you need to do regularly. This chapter provides a rudimentary overview of the inner workings of shell scripts and how they can be used. You learn how shell scripts are responsible for the messages that scroll by on the system console during booting and how simple scripts can be harnessed to a scheduling facility (such as `cron` or `at`) to simplify administrative tasks.

You also learn to fine-tune your machine to start at the most appropriate run level and to run only services you need. With that understanding, you'll be able to personalize your computer and cut down on the amount of time you spend repetitively typing the same commands.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Understanding Shell Scripts

Have you ever had a task that you needed to do over and over that took a lot of typing on the command line? Do you ever think to yourself, "Wow, I wish there was just one command I could type to do all this of this"? Maybe a shell script is what you're after.

Shell scripts are the equivalent of batch files in MS-DOS, and can contain long lists of commands, complex flow control, arithmetic evaluations, user-defined variables, user-defined functions, and sophisticated condition testing. Shell scripts are capable of handling everything from simple one-line commands to something as complex as starting up your Red Hat Linux system.
In fact, as you will read in this chapter, Red Hat Linux does just that. It uses shell scripts (`/etc/rc.d/rc.sysint` and `/etc/rc`) to check and mount all your filesystems, set up your consoles, configure your network, launch all your system services, and eventually provide you with your login screen. While there are nearly a dozen different shells available in Red Hat Linux, the default shell is called `bash`, the Bourne-Again shell.

Executing and debugging shell scripts

One of the primary advantages of shell scripts is that they can be opened in any text editor to see what they do. A big disadvantage is that shell scripts often execute more slowly than compiled programs. There are two basic ways to execute a shell script:

The filename is used as an argument to the shell (as in `bash myscript`). In this method, the file does not need to be executable; it just contains a list of shell commands. The shell specified on the command line is used to interpret the commands in the script file. This is most common for quick, simple tasks.
The shell script may also have the name of the interpreter placed in the first line of the script preceeded by a `#!` (as in `#!/bin/bash`), and have its execute bit set (using `chmod +x`). You can then run your script just like any other program in your path simply by typing the name of the script on the command line.

Cross-Reference See Chapter 4 for more on chmod and read/write/execute permissions.

When scripts are executed in either manner, options to the program may be specified on the command line. Anything following the name of the script is referred to as a *command-line argument*.

As with writing any software, there is no substitute to clear and thoughtful design and lots of comments. The pound sign (#) prefaces comments and can take up an entire line or exist on the same line as script code. It's best to implement more complex shell scripts in stages, making sure the logic is sound at each step before continuing. Here are a few ways to make sure things are working as expected during testing:

Place an `echo` statement at the beginning of lines within the body of a loop. That way, rather than executing the code, you can see what will be executed without making any permanent changes.
To achieve the same goal, you could place dummy `echo` statements throughout the code. If these lines get printed, you know the correct logic branch is being taken.
You could use `set +x` near the beginning of the script to display each command that is executed or launch your scripts using `sh -x` *`myscript`*.

Understanding shell variables

Often within a shell script, you want to reuse certain items of information. During the course of processing the shell script, the name or number representing this information may change. To store information used by a shell script in a way that it can be easily reused, you can set variables. Variable names within shell scripts are case-sensitive and can be defined in the following manner:

NAME=*value*

The first part of a variable is the variable name, and the second part is the value set for that name. Variables can be assigned from constants, like text or numbers. This is useful for initializing values or saving lots of typing for long constants. Here are examples where variables are set to a string of

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

characters (CITY) and a numeric value (PI):

CITY="Springfield"

PI=3.14159265

Variables can contain the output of a command or command sequence. You can accomplish this by either enclosing the command in backticks (`) or by enclosing the command in parentheses. This is a great way to get information that can change from computer to computer or from day to day. Here we set the output of the `uname -n` command to the MACHINE variable. Then we use parens to set NUM_FILES to the number of files in the current directory by piping ( | ) the output of the `ls` command to the word count command (`wc -l`).

MACHINE=`uname –n`

NUM_FILES=(/bin/ls | wc –l)

Variables can also contain the value of other variables. This is useful when you have to preserve a value that will change so you can use it later in the script. Here BALANCE is set to the value of the CurBalance variable.

BALANCE=$CurBalance

Note When assigning variables, use only the variable name (for example, BALANCE). When referenced, meaning you want the *value* of the variable, precede it with a dollar sign (as in $CurBalance).

Special Shell Variables

There are special variables that the shell assigns for you . The most commonly used are called the *positional parameters* or *command line arguments* and are referenced as $0, $1, $2, $3…$n. $0 is special and is assigned the name used to invoke your script; the remainder are assigned the values of the parameters passed on the command line . For instance, if the shell script named `myscript` were called as:

myscript foo bar

the positional parameter $0 would be `myscript`, `$1` would be `foo`, and `$2` would be `bar`.
Another variable, `$#`, tells you how many parameters your script was given. In our example, `$#` would be 2. Another particularly useful special shell variable is `$?`, which receives the exit status of the last command executed. Typically, a value of zero means everything is okay, and anything other than zero indicates an error of some kind. For a complete list of special shell variables, refer to the `bash` man page .

Parameter expansion in bash

As mentioned earlier, if you want the value of a variable, you precede it with a `$` (for example, `$CITY`). This is really just a shorthand for the notation `${CITY}`. Bash has special rules that allow you to expand the value of a variable in different ways. Going into all the rules is probably a little overboard for a quick introduction to shell scripts, but Table 12-1 presents some common constructs that you're likely to see in bash scripts you find on your Red Hat Linux box.

Table 12-1: Examples of bash parameter expansion

Construction	Meaning
`${var:-value}`	If variable is unset or empty, expand this to value
`${var#pattern}`	Chop the shortest match for pattern from the end of var's value
`${var##pattern}`	Chop the longest match for pattern from the end of var's value
`${var%pattern}`	Chop the shortest match for pattern from the front of var's value

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 12-1: Examples of bash parameter expansion

Construction	Meaning
`${var%%pattern}`	Chip the longest match for pattern from the front of var's value

Try typing the following commands from a shell to test out how parameter expansion works:

THIS="Example"

THIS=${THIS:-"Not Set"}

THAT=${THAT:-"Not Set"}

echo $THIS

Example

echo $THAT

Not Set

In the previous examples, the THIS variable is set to the word Example. In the next two lines, the THIS and THAT variables are set to their current values or to Not Set, if they are not currently set. Notice that because we just set THIS to the string Example, when we echo the value of THIS it appears as Example. However, since THAT was not set, it appears as Not Set.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note For the rest of this section, we show how variables and commands may appear in a shell script. To try out any of those examples, however, you can simply type them into a

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

In the following example, we set MYFILENAME to `/home/digby/myfile.txt`. Next, the FILE variable is set to `myfile.txt` and DIR is set to `/home/digby`. In the NAME variable, the file name is cut down to simply `myfile`, then in the EXTENSION variable the file extension is set to `txt`. (To try these out, you can type them at a shell prompt as we did with the previous example, then echo the value of each variable to see how it is set.)

```
MYFILENAME="/home/digby/myfile.txt"
FILE=${MYFILENAME##*/}      #FILE becomes "myfile.txt"
DIR=${MYFILENAME%/*}        #DIR becomes "/home/digby
NAME=${FILE%.*}             #NAME becomes "myfile"
EXTENSION=${FILE#*.}        #EXTENSION becomes "txt"
```

Performing arithmetic in shell scripts

Bash used *untyped* variables, meaning it normally treats variables as strings or text, but can change them on the fly if you want it to. Unless you tell it otherwise with `declare`, your variables are just a bunch of letters to bash. But when you start trying to do arithmetic with them, bash will convert them to integers if it can. This makes it possible to do some fairly complex arithmetic in bash.

Integer arithmetic can be performed using the built-in `let` command or through the external `expr` or `bc` commands. After setting the variable `BIGNUM` value to `1024`, the three commands that follow would all store the value `64` in the `RESULT` variable:

```
BIGNUM=1024
let RESULT=$BIGNUM/16
RESULT=`expr $BIGNUM / 16`
RESULT=`echo "$BUGNUM / 16" | bc –l`
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note While most elements of shell scripts are relatively freeform (where whitespace, such as spaces or tabs, is insignificant), both let and expr are particular about spacing. The let command

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To see a complete list of the kinds of arithmetic you can perform using the `let` command, type **help let** at the bash prompt.

Using programming constructs in shell scripts

One of the features that make shell scripts so powerful is that their implementation of looping and conditional execution constructs similar to those found in more complex scripting and programming languages. You can use several different types of loops, depending on your needs.

The "if… then" statements

The most commonly used programming construct is conditional execution, or the "if" statement. It is used to perform actions only under certain conditions. There are several variations, depending on whether you're testing one thing, or want to do one thing if a condition is true, but another thing if a condition is false, or if you want to test several things one after the other.
The first `if/then` example tests if VARIABLE is set to the number 1. If it is, then the `echo` command is used to say that it is set to 1.

```
if [ $VARIABLE -eq 1 ] ; then
echo "The variable is 1"
fi
```

Instead of using `-eq`, you can use the equals sign (=), as shown in the following example. Using the `else` statement, different words can be echoed if the criterion of the `if` statement isn't met (`$STRING = "Friday"`).

```
if [ $STRING = "Friday" ] ; then
echo "WhooHoo!  Friday!"
else
echo "Will Friday ever get here?"
fi
```

You can also reverse tests with an exclamation mark (!). In the following example, if STRING is not Monday, then "At least it's not Monday" is echoed.

```
if [ $STRING != "Monday" ] ; then
  echo "At least it's not Monday"
fi
```

In the following example, `elif` is used to test for an additional condition (is `filename` a file or a directory).

```
if [ -f $filename ] ; then
  echo "$filename is a regular file"
elif [ -d $filename ] ; then
  echo "$filename is a directory"
else
  echo "I have no idea what $filename is"
fi
```

As you can see from the preceding examples, the condition you are testing is placed between square brackets [ ]. When a test expression is evaluated, it will return either a value of 0, meaning that it is true, or a 1, meaning that it is false. Table 12-2 lists the conditions that are testable and is quite a handy reference. (If you're in a hurry, you can type **help test** on the command line to get the same information.)

Table 12-2: Operators for Test Expressions

Operator	What Is Being Tested?
`-a file`	Does the file exist? (same as `-e`)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 12-2: Operators for Test Expressions

Operator	What Is Being Tested?
`-b file`	Is the file a special block device?
`-c file`	Is the file character special (e.g., a character device)? Used to identify serial lines and terminal devices.
`-d file`	Is the file a directory?
`-e file`	Does the file exist? (same as `-a`)
`-f file`	Does the file exist, and is it a regular file (e.g., not a directory, socket, pipe, link, or device file)?
`-g file`	Does the file have the set-group-id bit set?
`-h file`	Is the file a symbolic link? (same as `-L`)
`-k file`	Does the file have the sticky bit set?
`-L file`	Is the file a symbolic link?
`-n string`	Is the length of the string greater than 0 bytes?
`-O file`	Do you own the file?
`-p file`	Is the file a named pipe?
`-r file`	Is the file readable by you?
`-s file`	Does the file exist, and is it larger than 0 bytes?
`-S file`	Does the file exist, and is it a socket?
`-t fd`	Is the file descriptor connected to a terminal?
`-u file`	Does the file have the set-user-id bit set?
`-w file`	Is the file writable by you?
`-x file`	Is the file executable by you?
`-z string`	Is the length of the string 0 (zero) bytes?
`expr1 -a expr2`	Are both the first expression and the second expression true?
`expr1 -o expr2`	Is either of the two expressions true?
`file1 -nt file2`	Is the first file newer than the second file (using the modification timestamp)?
`file1 -ot file2`	Is the first file older than the second file (using the modification timestamp)?
`file1 -ef file2`	Are the two files associated by a link (a hard link or a symbolic link)?

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 12-2: Operators for Test Expressions

Operator	What Is Being Tested?
`var1 = var2`	Is the first variable equal to the second variable?
`var1 -eq var2`	Is the first variable equal to the second variable?
`var1 -ge var2`	Is the first variable greater than or equal to the second variable?
`var1 -gt var2`	Is the first variable greater than the second variable?
`var1 -le var2`	Is the first variable less than or equal to the second variable?
`var1 -lt var2`	Is the first variable less than the second variable?
`var1 != var2`	Is the first variable not equal to the second variable?
`var1 -ne var2`	Is the first variable not equal to the second variable?

There is also a special shorthand method of performing tests that can be useful for simple *one-command* actions. In the following example, the two pipes (||) indicate that if the directory being tested for doesn't exist (`-d dirname`), then make the directory (`mkdir $dirname`).

```
# [ test ] || {action}
# Perform simple single command {action} if test is false
[ -d $dirname ] || mkdir $dirname
```

Instead of pipes, two ampersands can be used to test if something is true. In the example below, a command is being tested to see if it includes at least three command line arguments.

```
# [ test ] && {action}
# Perform simple single command {action} if test is true
[ $# -ge 3 ] && echo "There are at least 3 command line arguments."
```

The case command

Another frequently used construct is the `case` command. Similar to a `switch` statement in programming languages, this can take the place of several nested `if` statements. A general form of the `case` statement is as follows:

```
case "VAR" in
  Result1)
    { body };;
  Result2)
    { body };;
  *)
    { body } ;;
esac
```

One use for the `case` command might be to help with your backups. The following case statement tests for the first three letters of the current day (`case `date +%a` in`). Then, depending on the day, a particular backup directory (BACKUP) and tape drive (TAPE) is set.

```
# Our VAR doesn't have to be a variable,
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
# it can be the output of a command as well
# Perform action based on day of week
case `date +%a` in
  "Mon")
        BACKUP=/home/myproject/data0
        TAPE=/dev/rft0
# Note the use of the double semi-colon to end each option
        ;;
# Note the use of the "|" to mean "or"
  "Tue" | "Thu")
        BACKUP=/home/myproject/data1
        TAPE=/dev/rft1
        ;;
  "Wed" | "Fri")
        BACKUP=/home/myproject/data2
        TAPE=/dev/rft2
        ;;
# Don't do backups on the weekend.
  *)
        BACKUP="none"
        TAPE=/dev/null
        ;;
esac
```

The asterisk (`*`) is used as a catchall, similar to the `default` keyword in the C programming language. In this example, if none of the other entries are matched on the way down the loop, the asterisk is matched, and the value of BACKUP becomes `none`.

The "for . . . do" loop

Loops are used to perform actions over and over again until a condition is met or until all data has been processed. One of the most commonly used loops is the `for . . . do` loop. It iterates through a list of values, executing the body of the loop for each element in the list. The syntax and a few examples are presented here:

```
for VAR in LIST
do
   { body }
done
```

The `for` loop assigns the values in `LIST` to `VAR` one at a time. Then for each value, the body in braces between `do` and `done` is executed. `VAR` can be any variable name, and `LIST` can be composed of pretty much any list of values or anything that generates a list.

```
for NUMBER in 0 1 2 3 4 5 6 7 8 9
do
  echo The number is $NUMBER
done


for FILE in `/bin/ls`
do
  echo $FILE
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
done


# You can also write it this way, which is somewhat cleaner.
for NAME in John Paul Ringo George ; do
  echo $NAME is my favorite Beetle
done
```

Each element in the `LIST` is separated from the next by white space. This can cause trouble if you're not careful, since some commands, like `ls -l`, output multiple fields per line, each separated by whitespace.

If you're a die-hard C programmer, bash allows you to use C syntax to control your loops.

```
LIMIT=10
# Double parentheses, and no $ on LIMIT even though it's a variable!
for ((a=1; a <= LIMIT ; a++)) ; do
 echo  "$a"
done
```

The "while . . . do" and "until . . . do" loops

Two other possible looping constructs are the `while. . . do` loop and the `until. . . do` loop. The structure of each is presented here:

```
while condition      until condition
do                   do
  { body }             { body }
done                 done
```

Here is an example of a `while` loop that will output the number `0123456789`:

```
N=0
while [ $N –lt 10 ] ; do
  echo –n $N
  let N=$N+1
done
```

Another way to output the number 0123456789 is to use an until loop as follows:

```
N=0
until [ $N –eq 10 ] ; do
  echo –n $N
  let N=$N+1
done
```

Some useful external programs

Bash is great and has lots of built-in commands, but it usually needs some help to do anything really useful. Some of the most common useful programs you'll see used are `grep`, `cut`, `tr`, and `sed`. Like all the best UNIX tools, most of these programs are designed to work with standard input and standard output, so you can easily use them with pipes and shell scripts.

The General Regular Expression Parser (grep)

The name sounds intimidating, but `grep` is just a way to find patterns in files or text. Think of it as a useful search tool. Getting really good with regular expressions is quite a challenge, but many useful things can be accomplished with just the simplest forms.

For example, you can display a list of all regular user accounts by using `grep` to search for all lines that contain the text `/home` in the `/etc/passwd` file as follows:

```
grep /home /etc/passwd
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Or you could find all environment variables that begin with HO using the following command:

```
env | grep ^HO
```

To find a list of options to use with the grep command, type **man grep**.

Remove sections of lines of text (cut)

The `cut` command can extract specific fields from a line of text or from files. It is very useful for parsing system configuration files into easy-to-digest chunks. You can specify the field separator you want to use and the fields you want, or you can break a line up based on bytes.
The following example lists all home directories of users on your system. Using an earlier example of the `grep` command, this line pipes a list of regular users from the `/etc/passwd` file, then displays the sixth field (`-f6`) as delimited by a colon (`-d':'`).

```
grep /home /etc/passwd | cut –f6 –d':'
```

Translate or delete characters (tr)

The `tr` command is a character-based translator that can be used to replace one character or set of characters with another or to remove a character from a line of text.

The following example translates all upper-case letters to lowercase letters and displays the words "mixed upper and lower case" as a result:

```
FOO="Mixed UPpEr aNd LoWeR cAsE"

echo $FOO | tr [A-Z] [a-z]
```

In this example, the `tr` command is used on a list of file names to rename any files in that list so that any spaces contained in a file name are translated into underscores:

```
for file in * ; do
  d=`echo $file | tr [:blank:] [_]`
  [ "$file" –eq "-d" ] || mv "$file" "$d"
done
```

The Stream Editor (sed)

The `sed` command is a simple scriptable editor, and as such can perform only simple edits, such as removing lines that have text matching a certain pattern, replacing one pattern of characters with another, and other simple edits. To get a better idea of how `sed` scripts work, there's no substitute for the online documentation, but here are some examples of common uses.
You can use the `sed` command to essentially do what we did earlier with the `grep` example: search the `/etc/passwd` file for the word `home`. Here the sed command searches the entire `/etc/passwd` file, searches for the word home, and prints any line containing the word home.

```
sed –n –e '/home/p' /etc/passwd
```

In this example, `sed` searches the file `somefile.txt` and replaces every instance of the string `Mac` with `Linux`. The output is then sent to the `fixed_file.txt` file.

```
sed –e 's/Mac/Linux/' somefile.txt > fixed_file.txt
```

You can get the same result using a pipe:

```
cat somefile.txt | sed –e 's/Mac/Linux/' > fixed_file.txt
```

By searching for a pattern and replacing it with a null pattern, you delete the original pattern. This example searches the contents of the `somefile.txt` file and replaces extra blank spaces at the end of each line (`s/ *$`) with nothing (`//`). Results go to the `fixed_file.txt` file.

```
cat somefile.txt | sed –e 's/ *$//` > fixed_file.txt
```

Trying some simple shell scripts

Sometimes the simplest of scripts can be the most useful. If you type the same sequence of commands repetitively, it makes sense to store those commands (once!) in a file. Here are a couple of simple, but useful, shell scripts.

A simple telephone list

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

This idea has been handed down from generation to generation of old UNIX hacks. It's really quite simple, but it employs several of the concepts just introduced.

```
#!/bin/bash
# (@)/ph
# A very simple telephone list
# Type "ph new name number" to add to the list, or
# just type "ph name" to get a phone number

PHONELIST=~/.phonelist.txt

# If no command line parameters ($#), there
# is a problem, so ask what they're talking about.
if [ $# -lt 1 ] ; then
  echo "Whose phone number did you want?"
  exit 1
fi

# Did you want to add a new phone number?
if [ "$1" = "new" ] ; then
  shift
  echo $* >> $PHONELIST
  echo $* added to database
  exit 0
fi

# Nope. But does the file have anyting in it yet?
# This might be our first time using it, after all.
if [ ! -s $PHONELIST ] ; then
  echo "No names in the phone list yet!"
  exit 1
else
  grep –q "$*" $PHONELIST     # Quietly search the file
  if [ $? –ne 0 ] ; then          # Did we find anything?
    echo "Sorry, that name was not found in the phone list"
    exit 1
  else
    grep "$*" $PHONELIST
  fi
fi
exit 0
```

A Simple Backup Script

Since nothing works forever and mistakes happen, backups are just a fact of life when dealing with computer data. This simple script will back up all the data in the home directories of all the users on

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

your Red Hat Linux system.

```
#!/bin/bash
# (@)/my_backup
# A very simple backup script
#

TAPE=/dev/rft0

# Rewind the tape device $TAPE
mt $TAPE rew
# Get a list of home directories
HOMES=`grep /home /etc/passwd | cut –f6 –d';'`
# Backup the data in those directories
tar cvf $TAPE $HOMES
# Rewind and eject the tape.
mt $TAPE rewoffl
```

Cross-Reference See Chapter 13 for details on backing up and restoring files.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

System Start-Up and Shutdown

During system start-up, a series of scripts are run to start the services that you need. These include scripts to start network interfaces, mount directories, and monitor your system. Most of these scripts are run from subdirectories of `/etc/rc.d`. The program that starts most of these services up when you boot and stops them when you shut down is the `/etc/rc.d/rc` script. The following sections describe run-level scripts and what you can do with them.

Starting run-level scripts

As previously mentioned, the `/etc/rc.d/rc` script is a script that is integral to the concept of run levels. Any change of run level causes the script to be executed, with the new run level as an argument. Here's a quick run-down of what the `/etc/rc.d/rc` script does:

Checks that run level scripts are correct. The rc script checks to find each run level script that exists and excludes those that represent backup scripts left by rpm updates.

Determines current and previous run levels. Determines the current and previous run levels to know which run-level scripts to stop (previous level) and start (current level).

Decides whether to enter interactive startup. If the confirm option is passed to the boot loader at boot time, all server processes must be confirmed at the system console before starting.

Kills and starts run-level scripts. Stops run-level scripts from the previous level, then starts run-level scripts from the current level.

In Red Hat Linux, most of the services that are provided to users and computers on the network are started from run-level scripts.

Understanding run-level scripts

A software package that has a service to start at boot time (or when the system changes run levels) can add a script to the `/etc/init.d` directory. That script can then be linked to an appropriate run-level directory and either be started or stopped (to start or stop the service).

Table 12-5 lists many of the typical run-level scripts that are found in `/etc/init.d` and explains their function. Depending on the Red Hat Linux software packages you installed on your system, you may have dozens more run-level scripts than you see here. (Later, I describe how these files are linked into particular run-level directories.)

Table 12-5: Run-Level Scripts Contained in /etc/init.d

Run-Level Scripts	What Does It Do?
`apmd`	Controls the Advanced Power Management daemon, which monitors battery status, and which can safely suspend or shut down all or part of a machine that supports it.
`atd`	Starts or stops the at daemon to receive, queue, and run jobs submitted via the `at` or `batch` commands. (The `anacron` run-level script runs at and batch jobs that were not run because the computer was down.)
`autofs`	Starts and stops the `automount` daemon, for automatically mounting file systems (so, for example, a CD can be automatically mounted when it is inserted).
`crond`	Starts or stops the `cron` daemon to periodically run routine commands.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 12-5: Run-Level Scripts Contained in /etc/init.d

Run-Level Scripts	What Does It Do?
`cups-lpd`	Controls the printer daemon that handles spooling printing requests.
`dhcpd`	Starts or stops the `dhcpd` daemon, which automatically assigns IP addresses to computers on a LAN.
`gpm`	Controls the `gpm` daemon, which allows the mouse to interact with console- and text-based applications.
`halt`	Terminates all processes, writes out accounting records, removes swap space, unmounts all file systems, and either shuts down or reboots the machine (depending upon how the command was called).
`httpd`	Starts the `httpd` daemon, which allows your computer to act as an HTTP server (i.e., to serve Web pages).
`iptables`	Starts the `iptables` firewall daemon, which manages any `iptables`-style firewall rules set up for your computer.
`keytable`	Loads the predefined keyboard map.
`killall`	Shuts down any subsystems that may still be running prior to a shutdown or reboot.
`kudzu`	Detects and configures new hardware at boot time.
`netfs`	Mounts or unmounts network (NFS, SMB, and NCP) file systems.
`network`	Starts or stops all configured network interfaces and initializes the TCP/IP and IPX protocols.
`nfs`	Starts or stops the NFS-related daemons (`rpc.nfsd`, `rpc.mountd`, `rpc.statd`, and `rcp.rquotad`) and exports shared file systems.
`pcmcia`	Loads or unloads modules, drivers, and programs (including the `cardmgr` daemon) to support PCMCIA cards (Ethernet adapters, modems, memory cards, etc.) in laptop computers.
`portmap`	Starts or stops the `portmap` daemon, which manages programs and protocols that utilize the Remote Procedure Call (RPC) mechanism.
`random`	Loads or saves the current state of the machine's random number generator's random seed to ensure more random randomness.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 12-5: Run-Level Scripts Contained in /etc/init.d

Run-Level Scripts	What Does It Do?
`routed`	Starts or stops the `routed` daemon, which controls dynamic-routing table updates via the Router Information Protocol (RIP).
`rwhod`	Starts or stops the `rwhod` daemon, which enables others on the network to obtain a list of all currently logged-in users.
`sendmail`	Controls the `sendmail` daemon, which handles incoming and outgoing SMTP (Simple Mail Transport Protocol) mail messages.
`single`	Terminates all running processes and enters run level 1 (single-user mode).
`smb`	Starts or stops the `smbd` and `nmbd` daemons for allowing access to Samba file and print services.
`snmpd`	Starts or stops the `snmpd` (Simple Network Management Protocol) daemon, which enables others to view machine-configuration information.
`squid`	Starts or stops the `squid` services, which enables proxy service to clients on your network.
`syslog`	Starts or stops the `klogd` and `syslogd` daemons that handle logging events from the kernel and other processes, respectively.
`xfs`	Starts or stops `xfs`, the X Window font server daemon.
`xinetd`	Sets the machine's host name, establishes network routes, and controls `xinetd`, the network services daemon which listens for incoming TCP/IP connections to the machine.
`ypbind`	Binds to an NIS (Network Information Service) master server (if NIS is configured), and starts or stops the ypbind process, which communicates with the master server.

Each script representing a service that you want to start or stop is linked to a file in each of the run-level directories. For each run level, a script beginning with `K` stops the service, whereas a script beginning with `S` starts the service.

The two digits following the `K` or `S` in the filename provide a mechanism to select the priority in which the programs are run. For example, `S12syslog` is run before `S90crond`. However, while humans can readily see that 85 is less than 110, the file `S110my_daemon` is run before `S85gpm`. This is because the "ASCII collating sequence orders the files," which simply means that one positional character is compared to another. Therefore, a script beginning with the characters S110 is executed between `S10network` and `S15netfs` in run level 3.

All of the programs within the `/etc/rc`*X*`.d` directories (where *X* is replaced by a run-level number) are symbolic links, usually to a file in `/etc/init.d`. The `/etc/rc`*X*`.d` directories include the following:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

`/etc/rc0.d`: Run level 0 directory
`/etc/rc1.d`: Run level 1 directory
`/etc/rc2.d`: Run level 2 directory
`/etc/rc3.d`: Run level 3 directory
`/etc/rc4.d`: Run level 4 directory
`/etc/rc5.d`: Run level 5 directory
`/etc/rc6.d`: Run level 6 directory

In this manner, `/etc/rc0.d/K05atd` , `/etc/rc1.d/K05atd` , `/etc/rc2.d/K05atd` , `/etc/rc3.d/S95atd` , `/etc/rc4.d/S95atd` , `/etc/rc5.d/S95atd` , and `/etc/rc6.d/K05atd` are all symbolic links to `/etc/init.d/atd`. Using this simple, consistent mechanism, you can customize which programs are started at boot time.

Understanding what startup scripts do

Despite all the complicated rc*X*s, Ss, and Ks, the form of each start-up script is really quite simple. Because they are in plain text, you can just open one with a text editor to take a look at what it does. For the most part, a run-level script can be run with a `start` option, a `stop` option, and possibly a `restart` option. For example, the following lines are part of the contents of the `smb` script that defines what happens when the script is run with different options to start or stop the Samba file and print service:

```
#!/bin/sh
#
# chkconfig: - 91 35
# description: Starts and stops the Samba smbd and nmbd daemons \
#              used to provide SMB network services.
        .
        .
        .
start() {
        KIND="SMB"
        echo -n $"Starting $KIND services: "
        daemon smbd $SMBDOPTIONS
        RETVAL=$?
        echo
        KIND="NMB"
        echo -n $"Starting $KIND services: "
        daemon nmbd $NMBDOPTIONS
        RETVAL2=$?
        echo
        [ $RETVAL -eq 0 -a $RETVAL2 -eq 0 ] && touch /var/lock/subsys/smb || \
           RETVAL=1
        return $RETVAL
}

stop() {
        KIND="SMB"
        echo -n $"Shutting down $KIND services: "
        killproc smbd
        RETVAL=$?
        echo
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
        KIND="NMB"
        echo -n $"Shutting down $KIND services: "
        killproc nmbd
        RETVAL2=$?
        [ $RETVAL -eq 0 -a $RETVAL2 -eq 0 ] && rm -f /var/lock/subsys/smb
        echo ""
        return $RETVAL
}

restart() {
        stop
        start
}
      .
      .
      .
```

To illustrate what this script essentially does, I skipped some of the beginning and end of the script (where it checked if the network was up and running and set some values). Here are the actions `smb` takes when it is run with `start` or `stop`:

start — This part of the script starts the `smbd` and `nmbd` servers when the script is run with the start option.

stop — When run with the `stop` option, the `/etc/init.d/smb` script stops the `smbd` and `nmbd` servers.

The restart option runs the script with a `stop` option followed by a `start` option. If you want to start the `smb` service yourself, you can type the following command (as root user):

```
# /etc/init.d/smb start
Starting SMB services:          [ OK ]
Starting NMB services:          [ OK ]
```

To stop the service, you could type the following command:

```
# /etc/init.d/smb stop
Shutting down SMB services:     [ OK ]
Shutting down NMB services:     [ OK ]
```

The `smb` run-level script is different from other run-level scripts in that it supports several other options than `start` and `stop`. For example, this script has options (not shown in the example) that allow you to reload the `smb.conf` configuration file (`reload`) and check the status of the service (`rhstatus`).

Changing run-level script behavior

Modifying the start-up behavior of any such script merely involves opening the file in a text editor. For example, the `atd` daemon queues jobs submitted from the `at` and `batch` commands. Jobs submitted via `batch` are executed only if the system load is below a particular value, which can be set with a command-line option to the `atd` command. The default value of 0.8 is based on the assumption that a single-processor machine with less than 80 percent CPU utilization could handle the additional load of the batch job. However, if you were to add another CPU to your machine, the default threshold value would be too low and the batch jobs would not be sufficiently restricted.

You can change the system load threshold value from 0.8 to 1.6 to accommodate the increased processing capacity. To do this, simply modify the following line (in the `start` section) of the `/etc/init.d/atd` script:

```
daemon /usr/sbin/atd
```

Replace it with this line, using the `-l` argument to specify the new minimum system load value:

```
daemon /usr/sbin/atd -l 1.6
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

After saving the file and exiting the editor, you can reboot the machine or just run any of the following three commands to begin using the new batch threshold value:

/etc/init.d/atd reload

/etc/init.d/atd restart

/etc/init.d/atd stop ; /etc/rc.d/init.d/atd start

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Always make a copy of a run-level script before you change it. Also, keep track of changes you make to run-level scripts before you upgrade the packages they come from. You ne

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

If you are uncomfortable editing start-up scripts and you simply want to add options to the daemon process run by the script, there maybe a way of entering these changes without editing the start-up script directly. Check the `/etc/sysconfig` directory and see if there is a file by the same name as the script you want to modify. If there is, that file probably provides values that you can set to pass options to the start-up script. Sysconfig files exist for `amd`, `arpwatch`, `dhcpd`, `kudzu`, `ntpd`, `samba`, `squid`, and others.

Reorganizing or removing run-level scripts

There are several ways to deal with removing programs from the system start-up directories, adding them to particular run levels, or changing when they are executed. From a Terminal window, you can use the `chkconfig` command. From a GUI, use the Service Configuration window.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution You should never remove the run-level file from the /etc/init.d directory. Because no scripts are run from the /etc/init.d directory automatically, it is okay to keep them there. Scripts in

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To reorganize or remove run-level scripts from the GUI, use the Service Configuration window. Either click System Settings → Services or log in as root user and type the following from a Terminal window:

serviceconf &

Figure 12-1 shows an example of the Service Configuration window.

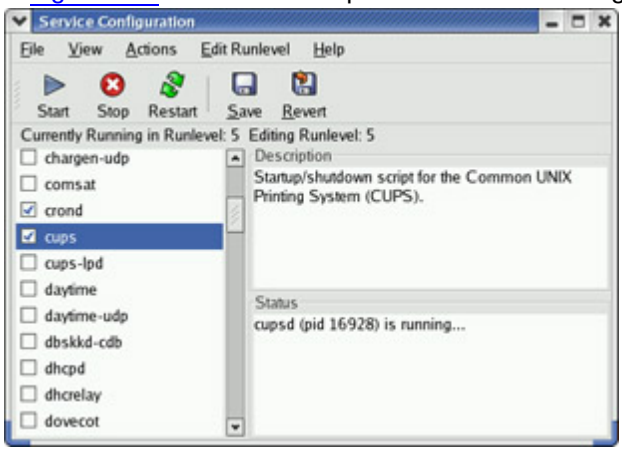

Figure 12-1: Reorganize, add, and remove run-level scripts from the Service Configuration window.

The Service Configuration window lets you reconfigure services for run levels 3, 4, and 5. The run levels that you are currently running and currently editing are displayed near the top of the screen. Services that are available for the run level appear in the middle of the frame, with check marks next to the ones configured to start at that level. Here is what you can do from this window:

Add — Click the box next to each service you want to start automatically at that run level so that a check mark appears in the box.

Remove — Click the run-level script that you want to remove for a particular run level to remove the check mark.

Save — Click File, then Save Changes on the window to save any changes you have made to the run-level scripts.

Refresh — Click View, then Refresh Service List to refresh the list of services.

Start, Stop, or Restart — Click a service on the list. Select Actions, then either Start, Stop, or Restart Service. The selected service immediately starts, stops, or restarts.

Some administrators prefer text-based commands for managing run-level scripts and for managing other system services that start automatically. The `chkconfig` command can be used to list whether services that run-level scripts start, as well as services the `xinetd` daemon starts, are on or off. To see a list of all system services, with indications that they are on or off, type the following:

chkconfig --list | less

You can then page through the list to see those services. If you want to view the status of an individual service, you can add the service at the end of the list option. For example, to see whether the nfs service starts in each run level, type the following:

chkconfig --list nfs

```
nfs     0:off   1:off   2:off   3:on    4:on    5:on    6:off
```

This example shows that the nfs service is set to be on for run levels 3, 4, and 5, but that it is set to off for run levels 0, 1, 2, and 6.

Another tool that can be run from the shell to change which services start and do not start at various levels is the `ntsysv` command. Type the following as root user from the shell:

ntsysv

A screen appears with a list of available services. Use the up and down arrow keys to locate the service you want. With the cursor on a service, press the Spacebar to toggle the service on or off. Press the Tab key to highlight the OK button, and press the Spacebar to save the change and exit. The `ntsysv` tool is somewhat clever; it checks your default run level (the initdefault set in the `/etc/inittab` file) and turns on the service so it starts at that run level.

Adding run-level scripts

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Suppose you want to create and configure your own run-level script. For example, after installing the binaries for the fictitious `my_daemon` program, it needs to be configured to start up in run levels 3, 4, and 5, and terminated in any other run level. You can add the script to the `/etc/init.d` directory, then use the `chkconfig` command to configure it.
To use `chkconfig`, ensure that the following lines are included in the `/etc/init.d/my_daemon` script:

```
# chkconfig: 345 82 28
# description: Does something pretty cool - you really \
#    have to see it to believe it!
# processname: my_daemon
```

With those lines in place, simply run the following command:

```
# chkconfig --add my_daemon
```

Appropriate links are created automatically. This can be verified with the following command:

```
# chkconfig --list my_daemon
```

The resulting output should look like:

```
my_daemon 0:off 1:off 2:off 3:on 4:on 5:on 6:off
```

Managing xinetd services
There are a bunch of services, particularly Internet services, that are not handled by separate run-level scripts. Instead, a single run-level script called `xinetd` (formerly `inetd`) is run to handle incoming requests for these services. For that reason, `xinetd` is sometimes referred to as the *super-server*. The `xinetd` run-level script (along with the `xinetd` daemon that it runs) offers the following advantages:

- **Fewer daemon processes.** Instead of one (or more) daemon processes running on your computer to monitor incoming requests for each service, the `xinetd` daemon can listen for requests for many different services. As a result, when you type `ps -ax` to see what processes are running, dozens of fewer daemon processes will be running than there would be if each service had its own daemon.
- **Access control and logging.** By using `xinetd` to oversee the management of services, consistent methods of access control (such as PAM) and consistent logging methods (such as the `/var/log/messages` file) can be used across all of the services.

When a request comes into your computer for a service that `xinetd` is monitoring, `xinetd` uses the `/etc/xinetd.conf` file to read configuration files contained in the `/etc/xinetd.d` directory. Then, based on the contents of the `xinetd.d` file for the requested service, a server program is launched to handle the service request (provided that the service is not disabled).
Each server process is one of two types: single-thread or multithread. A single-thread server will handle only the current request, whereas a multithread server will handle all incoming requests for the service as long as there is still a client holding the process open. Then the multithread server will close and `xinetd` will begin monitoring that service again.
The following are a few examples of services that are monitored by `xinetd`. The daemon process that is started up to handle each service is also listed.

- **comsat** (`/usr/sbin/in.comsat`) — Alerts the biff client that new mail has arrived.
- **eklogin** (`/usr/kerberos/sbin/klogind`) — Kerberos-related login daemon.
- **finger** (`/usr/sbin/in.fingerd`) — Handles incoming finger requests for information from remote users about local users.
- **gssftp** (`/usr/kerberos/sbin/ftpd`) — Kerberos-related daemon for handling file transfer requests (FTP).
- **imap** (`/usr/sbin/imapd`) — Daemon for handling requests from remote mail clients to get their mail from this IMAP server.
- **ntalk** (`/usr/sbin/in.ntalkd`) — Daemon for handling requests to set up chats between a remote user and a local one (using the `talk` command).
- **rlogin** (`/usr/sbin/in.rlogind`) — Daemon for responding to remote login requests (from a remote `rlogin` command).
- **rsh** (`/usr/sbin/in.rshd`) — Handles requests from a remote client to run a

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

command on the local computer.

Other services that can be launched by requests that come to `xinetd` include services for POP mail, remote telnet requests, Samba configuration requests (swat), and amanda network backups. A short description of each service is included in its `/etc/xinetd.d` file.

Manipulating run levels

Aside from the run level chosen at boot time (usually 3 or 5) and the shutdown or reboot levels (0 and 6, respectively), you can change the run level at any time while you're logged in (as root user). The `telinit` command (really just a symbolic link to `init`) enables you to specify a desired run level, causing the termination of all system processes that shouldn't exist in that run level, and starting all processes that should be running.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The telinit command is also used to instruct init to reload its configuration file, /etc/initttab. This is accomplished with either the telinit q or the telinit Q commands.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

For example, if you encountered a problem with your hard disk on startup, you may be placed in single-user mode (run level 1) to perform system maintenance. After the machine is stable, you can just execute the command as follows:

telinit 5

The `init` command handles terminating and starting all processes necessary to present you with a graphical login window.

Determining the current run level

You can determine the machine's current run level with the aptly named `runlevel` command. Using the previous example of booting into single-user mode and then manually changing the run level, the output of the `runlevel` command would be:

runlevel

S 5

This means that the previous run level was S (for single-user mode) and the current run level is 5. If the machine had booted properly, the previous run level would be listed as N to indicate that there really wasn't a previous run level.

Changing to a shutdown run level

Shutting down the machine is simply a change in run level. With that in mind, other ways to change the run level include the `reboot`, `halt`, `poweroff`, and `shutdown` commands. The `reboot` command, which is a symbolic link to the `halt` command, executes a `shutdown -r` now, terminating all processes and rebooting the machine. The `halt` command executes `shutdown -h now`, terminating all processes and leaving the machine in an idle state (but still powered on). Similarly, the `poweroff` command, which is also a link to the `halt` command, executes a change to run level 0, but if the machine's BIOS supports Advanced Power Management (APM), it will switch off the power to the machine.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note A time must be given to the shutdown command, either specified as +m (representing the number of minutes to delay before beginning shutdown) or as hh:mm (an absolute time value, where

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Scheduling System Tasks

Frequently, you need to run a process unattended or at off-hours. The `at` facility is designed to run such jobs at specific times. Jobs you submit are spooled in the directory `/var/spool/at`, awaiting execution by the `at` daemon `atd`. The jobs are executed using the current directory and environment that was active when the job was submitted. Any output or error messages that haven't been redirected elsewhere are e-mailed to the user that submitted the job.

The following sections describe how to use the `at`, `batch`, and `cron` facilities to schedule tasks to run at specific times. These descriptions also include ways of viewing which tasks are scheduled and deleting scheduled tasks that you don't want to run anymore.

Using at.allow and at.deny

There are two access control files designed to limit which users can use the `at` facility. The file `/etc/at.allow` contains a list of users that are granted access, and the file `/etc/at.deny` contains a similar list of those who may not submit `at` jobs. If neither file exists, only the superuser is granted access to `at`. If a blank `/etc/at.deny` file exists (as in the default configuration), all users are allowed to utilize the `at` facility to run their own `at` jobs.

Specifying when jobs are run

There are many different ways to specify the time at which an `at` job should run (most of which look like spoken commands). Table 12-6 has a few examples. These are not complete commands — they only provide an example of how to specify the time that a job should run.

Table 12-6: Samples for Specifying Times in an at Job

Command Line	When the Command Is Run
`at now`	The job is run immediately.
`at now + 2 minutes`	The job will start 2 minutes from the current time.
`at now + 1 hour`	The job will start one hour from the current time.
`at now + 5 days`	The job will start five days from the current time.
`at now + 4 weeks`	The job will start four weeks from the current time.
`at now next minute`	The job will start in exactly 60 seconds.
`at now next hour`	The job will start in exactly 60 minutes.
`at now next day`	The job will start at the same time tomorrow.
`at now next month`	The job will start on the same day and at the same time next month.
`at now next year`	The job will start on the same date and at the same time next year.
`at now next fri`	The job will start at the same time next Friday.
`at teatime`	The job will run at 4 p.m. They keywords noon and midnight can also be used.
`at 16:00 today`	The job will run at 4 p.m. today.
`at 16:00 tomorrow`	The job will run at 4 p.m. tomorrow.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 12-6: Samples for Specifying Times in an at Job

Command Line	When the Command Is Run
`at 2:45pm`	The job will run at 2:45 p.m. on the current day.
`at 14:45`	The job will run at 2:45 p.m. on the current day.
`at 5:00 Sep 14 2003`	The job will begin at 5 a.m. on September 14, 2003.
`at 5:00 9/14/03`	The job will begin at 5 a.m. on September 14, 2003.

Submitting scheduled jobs

The `at` facility offers a lot of flexibility in how you can submit scheduled jobs. There are three ways to submit a job to the `at` facility:

Piped in from standard input. For example, the following command will attempt to build the Perl distribution from source in the early morning hours while the machine is likely to be less busy:

echo "cd /tmp/perl; make ; ls -al" | at 2am tomorrow

An ancillary benefit to this procedure is that a full log of the compilation process will be e-mailed to the user that submitted the job.

Read as standard input. If no command is specified, `at` will prompt you to enter commands at the special `at>` prompt, as shown in the following example. You must indicate the end of the commands by pressing Ctrl+D, which signals an End of Transmission (<EOT>) to at.

$ **at 23:40**
at> **cd /tmp/perl**
at> **make**
at> **ls -al**
at> ***<Ctrl-d>***

Read from a file. When the `-f` command-line option is followed by a valid filename, the contents of that file are used as the commands to be executed, as in the following example:

$ **at -f /root/bin/runme now + 5 hours**

This runs the commands stored in `/root/bin/runme` in five hours. The file can either be a simple list of commands or a shell script to be run in its own subshell (that is, the file begins with `#!/bin/bash` or the name of another shell).

Viewing scheduled jobs

You can use the `atq` command (effectively the same as `at -l`) to view a list of your pending jobs in the at queue, showing each job's sequence number, the date and time the job is scheduled to run, and the queue in which the job is being run.

The two most common queue names are "a," which represents the at queue, and "b," which represents the batch queue. All other letters (upper- and lowercase) can be used to specify queues with lower priority levels. If the `atq` command lists a queue name as =, it indicates that the job is currently running. Here is an example of output from the `atq` command:

atq

2 2003-09-02 00:51 a

3 2003-09-02 00:52 a

4 2003-09-05 23:52 a

Here you can see that there are three `at` jobs pending (job numbers 2, 3, and 4, all indicated as "a"). After the job number, the output shows the date and hour each job is scheduled to run.

Deleting scheduled jobs

If you decide that you'd like to cancel a particular job, you can use the `atrm` command (equivalent to `at -d`) with the job number (or more than one) as reported by the `atq` command. For example, using the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

following output from `atq`:

atq

18 2003-09-01 03:00 a

19 2003-09-29 05:27 a

20 2003-09-30 05:27 a

21 2003-09-14 00:01 a

22 2003-09-01 03:00 a

you can remove the jobs scheduled to run at 5:27 a.m. on September 29 and September 30 from the queue with the command `atrm 19 20`.

Using the batch command

If system resources are at a premium on your machine, or if the job you submit can run at a priority lower than normal, the `batch` command (equivalent to `at -q b`) may be useful. It is controlled by the same `atd` daemon, and it allows job submissions in the same format as `at` submissions (although the time specification is optional).

However, to prevent your job from usurping already scarce processing time, the job will run only if the system load average is below a particular value. The default value is 0.8, but specifying a command-line option to `atd` can modify this. This was used as an example in the earlier section describing start-up and shutdown. Here is an example of the `batch` command:

$ **batch**

at> **du -h /home > /tmp/duhome**

at> ***<Ctrl-d>***

In this example, after I type the `batch` command, the `at` facility is invoked to enable me to enter the command(s) I want run. Typing the `du -h /home > /tmp/duhome` command line has the disk usages for everything in the `/home` directory structure output to the `/tmp/duhome` file. On the next line, pressing Ctrl+D ends the batch job. As soon as the load average is low enough, the command is run. (Run the `top` command to view the current load average.)

Using the cron facility

Another way to run commands unattended is via the `cron` facility. Part of the vixie-cron RPM package, `cron` addresses the need to run commands periodically or routinely (at least, more often than you'd care to manually enter them) and allows lots of flexibility in automating the execution of the command. As with the `at` facility, any output or error messages that haven't been redirected elsewhere are e-mailed to the user that submitted the job.

Also like the `at` facility, `cron` includes two access control files designed to limit which users can use it. The file `/etc/cron.allow` contains a list of users that are granted access, and the file `/etc/cron.deny` contains a similar list of those who may not submit `cron` jobs. If neither file exists, all users are granted access to `cron`.

There are four places where a job can be submitted for execution by the cron daemon `crond`:

The `/var/spool/cron/`*`username`* file. This method, where each individual user (indicated by *username*) controls his or her own separate file, is the method used on UNIX System V systems.

The `/etc/crontab` file. This is referred to as the *system crontab file*, and was the original crontab file from BSD UNIX and its derivatives. Only root has permission to modify this file.

The `/etc/cron.d` directory. Files placed in this directory have the same format as the `/etc/crontab` file. Only root is permitted to create or modify files in this directory.

The `/etc/cron.hourly`, `/etc/cron.daily`, `/etc/cron.weekly`, and `/etc/cron.monthly` directories. Each file in these directories is a shell script that runs at the times specified in the `/etc/crontab` file (by default at one minute after the hour, at 4:02 a.m. every day, Sunday at 4:22 a.m., and 4:42 a.m. on the first day of the month, respectively). Only root is allowed to create or modify files in these directories.

The standard format of an entry in the `/var/spool/cron/`*`username`* file consists of five fields specifying when the command should run: minute, hour, day of the month, month, and day of the week.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The sixth field is the actual command to be run.
The files in the `/etc/cron.d` directory and the `/etc/crontab` file use the same first five fields to determine when the command should run. However, the sixth field represents the name of the user submitting the job (because it cannot be inferred by the name of the file as in a `/var/spool/cron/`*`username`* directory), and the seventh field is the command to be run. Table 12-7 lists the valid values for each field common to both types of files.

Table 12-7: Valid /etc/crontab Field Values

Field Number	Field	Acceptable Values
1	minute	Any integer between 0 and 59
2	hour	Any integer between 0 and 23
3	day of the month	Any integer between 0 and 31
4	month	Any integer between 0 and 12, or an abbreviation for the name of the month (Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec)
5	day of the week	Any integer between 0 and 7 (where both 0 and 7 can represent Sunday), or abbreviation for the day (Sun, Mon, Tue, Wed, Thu, Fri, Sat)

An asterisk (*) in any field indicates all possible values for that field. For example, an asterisk in the second column is equivalent to 0,1,2 . . . 22,23, and an asterisk in the fourth column means Jan,Feb,Mar . . . Nov,Dec. In addition, lists of values, ranges of values, and increments can be used. For example, to specify the days Monday, Wednesday, and Friday, the fifth field could be represented as the list Mon,Wed,Fri. To represent the normal working hours in a day, the range 9–5 could be specified in the second field. Another option is to use an increment, as in specifying 0–31/3 in the third field to represent every third day of the month, or */5 in the first field to denote every five minutes. Lines beginning with a # character in any of the crontab-format files are comments, which can be very helpful in explaining what task each command is designed to perform. It is also possible to specify environment variables (in Bourne shell syntax, e.g., `NAME="value"`) within the `crontab` file. Any variable can be specified to fine-tune the environment in which the job will run, but one that may be particularly useful is `MAILTO`. The following line will send the results of the `cron` job to a user other than the one that submitted the job:

```
MAILTO=otheruser
```

If the following line appears in a `crontab` file, all output and error messages that haven't already been redirected will be discarded:

```
MAILTO=
```

Modifying scheduled tasks with crontab

The files in `/var/spool/cron` should not be edited directly. They should only be accessed via the `crontab` command. To list the current contents of your own personal `crontab` file, type the following command:

$ **crontab -l**

All `crontab` entries can be removed with the following command:

$ **crontab -r**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Even if your personal `crontab` file doesn't exist, you can use the following command to begin editing it:

$ **crontab -e**

The file automatically opens in the text editor that is defined in your `EDITOR` or `VISUAL` environment variables, with vi as the default. When you're done, simply exit the editor. Provided there were no syntax errors, your `crontab` file will be installed. For example, if your user name is `jsmith`, you have just created the file `/var/spool/cron/jsmith`. If you add a line (with a descriptive comment, of course) to remove any old core files from your source code directories, that file may look similar to this:

```
# Find and remove core files from /home/jsmith/src
5 1 * * Sun,Wed find /home/jsmith/src -name core -exec rm {} \; > /dev/null 2>&1
```

The root user can access any user's individual `crontab` file by using the `-u` *`username`* option to the `crontab` command.

Understanding cron files

There are separate `cron` directories set up to contain `cron` jobs that run hourly, daily, weekly, and monthly. These `cron` jobs are all set up to run from the `/etc/crontab` file. The default `/etc/crontab` file looks like this:

```
SHELL=/bin/bash
PATH=/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root
HOME=/

# run-parts
01 * * * * root run-parts /etc/cron.hourly
02 4 * * * root run-parts /etc/cron.daily
22 4 * * 0 root run-parts /etc/cron.weekly
42 4 1 * * root run-parts /etc/cron.monthly
```

The first four lines initialize the run-time environment for all subsequent jobs (the subshell in which jobs will run, the executable program search path, the recipient of output and error messages, and that user's home directory). The next five lines execute (as the user root) the `run-parts` program that controls programs that you may wish to run periodically.

`run-parts` is a shell script that takes a directory as a command-line argument. It then sequentially runs every program within that directory (shell scripts are most common, but binary executables and links are also evaluated). The default configuration executes programs in `/etc/cron.hourly` at one minute after every hour of every day; `/etc/cron.daily` at 4:02 a.m. every day; `/etc/cron.weekly` at 4:22 a.m. on Sundays; and `/etc/cron.monthly` at 4:42 a.m. on the first day of each month.

Here are examples of files that are installed in `cron` directories for different software packages:

- **`/etc/cron.daily/logrotate.cron`** — Automates rotating, compressing, and manipulating system logfiles.
- **`/etc/cron.daily/makewhatis.cron`** — Updates the whatis database (contains descriptions of man pages), which is used by the `man -k`, `apropos`, and `whatis` commands to find man pages related to a particular word.
- **`/etc/cron.daily/slocate.cron`** — Updates the `/var/lib/slocate/slocate.db` database (using the `updatedb` command), which contains a searchable list of files on the machine.
- **`/etc/cron.daily/tmpwatch`** — Removes files from `/tmp`, `/var/tmp`, and `/var/catman` that haven't been accessed in ten days.
- **`/etc/cron.hourly/diskcheck`** — Checks available disk space on your hard drive each hour. Responds to low disk space based on settings in `/etc/diskcheck.conf`.

The `makewhatis.cron` script installed in `/etc/cron.weekly` is similar to the one in `/etc/cron.daily`, but it completely rebuilds the `whatis` database, rather than just updating the existing database.

Finally, in the `/etc/cron.d` directory are files that have the same format as `/etc/crontab` files.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

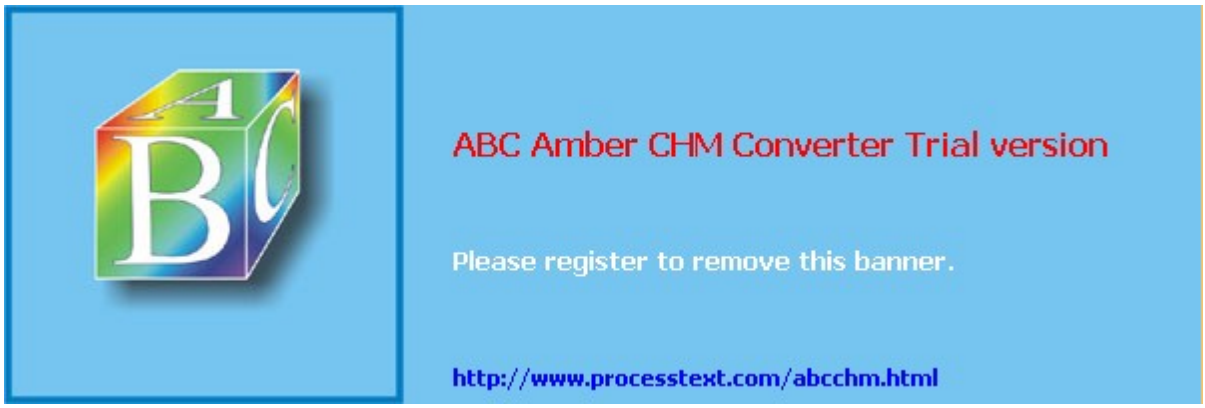

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

Shell scripts are an integral part of the Red Hat Linux system for configuring, booting, administering, and customizing Red Hat Linux. They are used to eliminate typing repetitive commands. They are frequently executed from the scheduling facilities within Red Hat Linux, allowing much flexibility in determining when and how often a process should run. And they control the start-up of most daemons and server processes at boot time.

The init daemon and its configuration file, `/etc/inittab`, also factor heavily in the initial start-up of your Red Hat Linux system. They implement the concept of run levels that is carried out by the shell scripts in `/etc/rc.d/init.d`, and they provide a means by which the machine can be shut down or rebooted in an orderly manner.

To have shell scripts configured to run on an ongoing basis, you can use the `cron` facility. `Cron` jobs can be added by editing `cron` files directly or by running commands such as `at` and `batch` to enter the commands to be run.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Chapter 13: Backing Up and Restoring Files

Overview

If you've ever suffered a hard drive crash, you know just how aggravating it can be. Irreplaceable data can be lost. Countless hours may be spent reinstalling your operating system and applications. It is not a fun experience. It need happen only once for you to learn the importance of making regular backups of your critical data.

Today, larger and faster backup media can simplify the process of backing up your data. Red Hat Linux supports many different types of media — such as writable CD-ROM, DVD, and magnetic tape — for creating backups. Using tools such as `cron`, you can configure backups to run unattended at scheduled times.

This chapter describes how to create a backup strategy and how to select media for backing up data on your Red Hat Linux system. It tells you how to do automated backups and backups over a network. It also describes how to restore individual files, or entire file systems, using tools such as the `restore` command.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Selecting a Backup Strategy

There are several approaches to backing up your data. You need to ask yourself a few questions to decide which approach is best for you. Some things that you should consider are:

- In the event of a crash, how much downtime can I tolerate?
- Will I need to recover older versions of my files or is the most recent revision sufficient?
- Do I need to back up files for just one computer or for many computers on a network?

Your answers to these questions will help you decide how often to do full backups and how often to do incremental backups. If the data is particularly critical, you may even decide that you need to have your data duplicated constantly, using a technique called *disk mirroring*. The following sections describe different backup methods.

Full backup

A full backup is one that stores every file on a particular disk or partition. If that disk should ever crash, you can rebuild your system by restoring the entire backup to a new disk. Whatever backup strategy you decide on, some sort of full backup should be part of it. You may perform full backups every night or perhaps only once every week; it depends on how often you add or modify files on your system, as well as the capacity of your backup equipment.

Incremental backup

An incremental backup is one that contains only those files that have been added or modified since the last time a more complete backup was made. You may choose to do incremental backups to conserve your backup media. Incremental backups also take less time to complete. This can be important when systems are in high use during the work week and running a full backup would degrade system performance. Full backups can be reserved for the weekend when the system is not in use.

Disk mirroring

Full and incremental backups can take time to restore, and sometimes you just can't afford that downtime. By duplicating your operating system and data on an additional hard drive, you can greatly increase the speed with which you can recover from a server crash.

With disk mirroring, it is usually common for the system to continuously update the duplicate drive with the most current information. In fact, with a type of mirroring called RAID 1 (described in Chapter 10), the duplicate drive is written to at the same time as the original, and if the main drive fails, the duplicate can immediately take over. This is called *fault-tolerant* behavior, which is a must if you are running a mission-critical server of some kind.

Network backup

All of the preceding backup strategies can be performed over a network. This is good because you can share a single backup device with many computers on a network. This is much cheaper and more convenient than installing a tape drive or other backup device in every system on your network. If you have many computers, however, your backup device will require a lot of capacity. In such a case, you may consider a mechanical tape loader, writable DVD drive or CD jukebox.

It is even possible to do a form of disk mirroring over the network. For example, a Web server may store a duplicate copy of its data on another server. If the first server crashes, a simple TCP/IP host name change can redirect the Web traffic to the second server. When the original server is rebuilt, it can recover all of its data from the backup server and be back in business.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Backing Up to a Hard Drive

Removable media such as tapes, DVDs, and CDs are not the only choice for backing up your data. You may find it useful to install a second hard drive in your system and use that drive for backups. This has several advantages over other backup media:

- Data can be backed up quickly and throughout the day; thus, backed-up data will be more current in the event of a crash.
- There's no medium to load. Data can be located and recovered more quickly.
- You can configure the second disk to be a virtual clone of the first one. If the first disk crashes, you can simply boot off of the second disk rather than installing new hardware. With disk mirroring software, this process can even be automated.

There are, however, a few disadvantages to backing up to a hard drive:

- This method is not well suited to keeping historical archives of the many revisions of your files because the hard drive will eventually fill up.
- The backups cannot be stored offsite, so if the server is destroyed, your data is lost.

Because of these disadvantages, you should probably not use a hard drive as your only backup medium. Instead, combine it with some sort of removable media backup such as a tape drive. Hard-drive mirroring can provide a *snapshot* backup that will get you up and running again quickly, while removable backups can provide longer-term storage of data and protection from more catastrophic failures.
Hard-drive mirroring is discussed in the next section.
The simplest form of second-hard-drive backup is to simply copy important files to the other drive using the `cp` or `tar` command. The most sophisticated method is to provide fault-tolerant disk mirroring using RAID software. If you don't need the high level of protection from data loss that RAID disk mirroring provides, but you still like the idea of having a complete duplicate of your data on-hand and ready for use, there is an alternative.

Getting and installing mirrordir to clone directories

The mirrordir package is a way of doing hard-drive mirroring. Mirrordir is a powerful tool that enables you to make and maintain an exact copy of a hierarchy of directories. You can find its official Web site at `http://mirrordir.sourceforge.net`. It can be downloaded by selecting the Download RPM button from that site.
After downloading the file, install it in the same way you install any rpm. For example, if you have downloaded the rpm file to `/tmp`, you can type:

```
# rpm -i /tmp/mirrordir*rpm
```

Cloning a directory with mirrordir

Now that you have `mirrordir` installed, you can use it to clone a directory. Suppose you have a second hard drive with a partition large enough to hold a copy of your `/home` partition. Your first step is to create a directory to mount the partition on, and then mount it. Log in as root and type the following:

```
# mkdir -p /mirror/home
```

```
# mount /dev/hdb5 /mirror/home
```

In this example, you are mounting the fifth partition (`5`) of the second hard drive (`hdb`), hence the device name `/dev/hdb5`. The `b` refers to the second disk, and the `5` refers to the partition number. You may use a different drive or partition. If so, adjust the device name accordingly. Assuming the `mount` command successfully mounted the drive, you are ready to copy your `/home` partition. Enter the following command:

```
# mirrordir /home /mirror/home
```

This command causes the contents of the `/home` directory to be exactly duplicated in the `/mirror/home` directory.

Caution It is very important that you get the order of these parameters correct. If you reverse them, the entire contents of your /home directory will be deleted! You will, in essence, be copying an empty directory back to your /home directory.

You now have a backup of your entire `/home` partition. You can run `mirrordir` again in the future, and it will again make the `/mirror/home` directory an exact duplicate of `/home`. It will copy to `/mirror/home` only those files that have been added or modified since the previous run of

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

`mirrordir`. Also, it will delete any files from `/mirror/home` that are no longer on `/home`. Thus, the mirror is kept current without copying the entire `/home` partition each time. If the disk with the `/home` partition should ever crash, you can replace the disk and then copy the file system back by issuing the `mirrordir` command with the parameters reversed:

mirrordir /mirror/home /home

Furthermore, you can be up and running even faster by simply turning the mirrored partition into the actual home partition. Just edit the `/etc/fstab` file and change the device name for `/home` so that it matches the mirrored directory. Now unmount and remount the `/home` partition (or reboot the computer), and the mirrored directory will be mounted to `/home`. Your users will never be the wiser. You may even consider mirroring every partition on your main disk. Then, even a complete disk crash can be recovered from quickly. Simply turn the secondary disk into the primary disk, and you are up and running again.

Automating mirroring

To automate the mirroring process, I suggest creating a small script that mounts the mirror partition, runs the `mirrordir` command, and then unmounts the partition. Leaving the partition unmounted when not in use reduces the chance of data being accidentally deleted from it. Create a script named `mirror.sh` and place it in the `/usr/local/sbin` directory. Something similar to the following should do the trick:

```
#!/bin/sh
#
#  mirror.sh:  Mirror the /home partition to a second hard drive
#
/bin/mount /dev/hdb5 /mirror/home
mirrordir /home /mirror/home
/bin/umount /mirror/home
```

Now tell `cron` to periodically run this script. Invoke the `cron` command with the `-e` option. This brings up an editor containing root's `cron` jobs. Add a line similar to the following at the end of the list.

```
0 * * * * /usr/local/sbin/mirror.sh
```

Save and exit the editor. This cron entry causes the `mirror.sh` script to run once an hour. If you decide to mirror other partitions, you can do it quite easily by creating appropriate mount points in `/mirror` and then adding matching sections to the `mirror.sh` script.

Team LiB

PREVIOUS NEXT

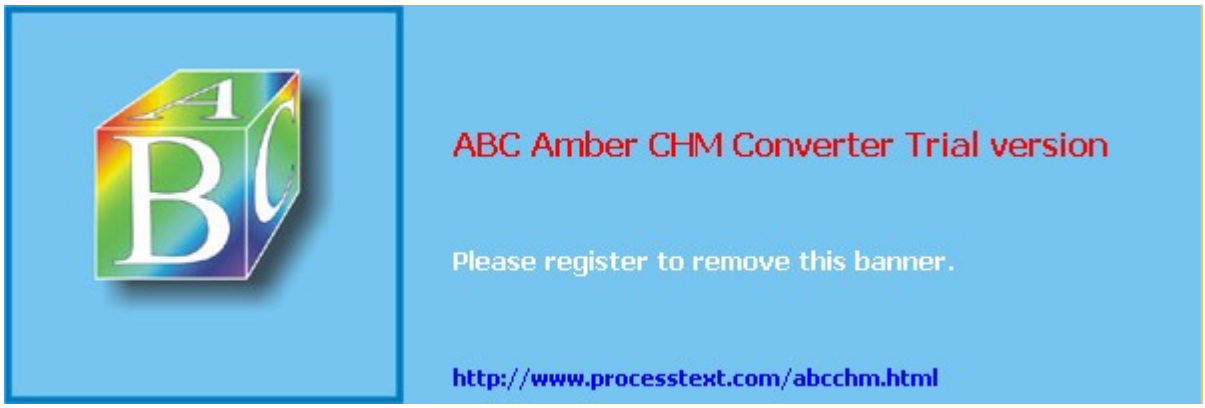

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Backing Up Files with dump

The `dump` command is the most commonly used tool for performing backups on UNIX systems. This command traces its history back to the early days of UNIX and thus is a standard part of nearly every version of UNIX. Likewise, the dump package is included in Red Hat Linux. If it was not installed by default when you first set up your Linux system, you can install it from the `dump` RPM file located on Red Hat Linux installation CD #1.

The dump package actually consists of several commands. You can read online man pages for more information about them, but Table 13-2 has a short description of the programs.

Table 13-2: Programs in the dump Package

Command	Description
`dump`	Creates backup archives of whole disk partitions or selected directories.
`restore`	Can be used to restore an entire archive or individual files from an archive to the hard drive.
`rmt`	A program used by the `dump` and `restore` commands to copy files across the network. You should never need to use this command directly.

Creating a backup with dump

When making a file system backup using the `dump` command, you must supply parameters specifying the dump level, the backup media, and the file system to back up. You can also supply optional parameters to specify the size of the backup media, the method for requesting the next tape, and the recording of file system dump times and status.

The first parameter to dump is always a list of single-letter option codes. This is followed by a space-separated list of any arguments needed by those options. The arguments appear in the same order as the options that require them. The final parameter is always the file system or directory being backed up.

dump *options arguments filesystem*

Table 13-3 lists the various one-letter option codes for the `dump` command.

Table 13-3: Options to dump

Dump Options	Description
`0-9`	The dump level. Selecting a dump level of `0` backs up all files (a full dump). A higher number backs up only those files modified since the last dump of an equal or lower number (in essence, an incremental dump). The default dump level is `9`.
`-B records`	The number of dump records per volume. Basically, the amount of data you can fit on a tape. This option takes a numeric argument.
`-b kbperdump`	The number of kilobytes per dump record. Useful in combination with the `-B` option. This option takes a numeric argument.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 13-3: Options to dump

Dump Options	Description
`-h level`	Files can be marked with a `nodump` attribute. This option specifies the dump level at or above which the `nodump` attribute is honored. This option takes a numeric argument of `1-9`.
`-f file`	The name of the file or device to write the dump to. This can even be a file or device on a remote machine.
`-d density`	Sets the tape density. The default is 1600 bits per inch. This option takes a numeric argument.
`-n`	When a dump needs attention (such as to change a tape), dump will send a message to all of the users in the operator group. This option takes no arguments.
`-s feet`	Specifies the length in feet of the dump tape. This calculation is dependent on tape density (option `d`) and the dump record (options `B` and `b`). This takes a numeric argument.
`-u`	Record this backup in the `/etc/dumpdates` file. It is a good idea to use this option, especially if you create incremental backups.
`-t date`	Specify a date and time on which to base incremental backups. Any files modified or added after that time will be backed up. This option causes `dump` to ignore the `/etc/dumpdates` file. It takes a single argument, a date in the format specified by the ctime man page.
`-W`	This option causes `dump` to list the file systems that need to be backed up. It does this by looking at the `/etc/dumpdates` file and the `/etc/fstab` file.
`-w`	This works like the `W` option but lists the individual files that should be backed up.

Thus, a typical `dump` command may look similar to the following:

dump 0uBf 500000 /dev/qft0 /dev/hda6

This command results in `dump` performing a level zero (full) backup of the `/dev/hda6` file system, storing the backup on the tape drive `/dev/qft0`, and recording the results in `/etc/dumpdates`. The `B` option is used to increase the expected tape block count to 500000; otherwise, `dump` would prompt for a new tape far earlier than required. The `dump` command prints status messages to the screen, letting you know how far along the backup has progressed and estimating how much time it will take to complete. The output looks similar to this:

DUMP: Date of this level 0 dump: Sat Aug 23 23:33:37 2003

DUMP: Dumping /dev/hda6 (/home) to /dev/qft0

DUMP: Exclude ext3 journal inode 8

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
DUMP: Label: /home
DUMP: mapping (Pass I) [regular files]
DUMP: mapping (Pass II) [directories]
DUMP: estimated 93303 tape blocks on 0.19 tape(s).
DUMP: Volume 1 started with block 1 at: Sat Aug 23 23:33:47
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
DUMP: Closing /dev/qft0
DUMP: Volume 1 completed at: Sat Aug 23 23:35:35 2003
DUMP: Volume 1 94360 tape blocks (92.15MB)
DUMP: Volume 1 took 0:01:48
DUMP: Volume 1 transfer rate: 873 kB/s
DUMP: 94360 tape blocks (92.15MB) on 1 volume(s)
DUMP: finished in 108 seconds, throughput 873 kBytes/sec
DUMP: Date of this level 0 dump: Sat Aug 23  23:33:37 2003
DUMP: Date this dump completed: Sat Aug 23  23:35:35 2003
DUMP: Average transfer rate: 873 kB/s
DUMP: DUMP IS DONE
```

Understanding dump levels

The `dump` command can back up all files on a file system, or it can selectively back up only those files that have changed recently. The dump level parameter is used to specify this behavior. A dump level of `0` results in a full backup of all files on the file system. Specifying a higher number (`1–9`) backs up only those files that have been changed or added since the most recent dump of the same or lower dump level. I recommend you use dump levels to implement a full and incremental backup schedule similar to that shown in Table 13-4.

Table 13-4: Recommended dump Schedule

Day of Week	Dump Level
Sunday	Level 0 (full dump). Eject the tape when done.
Monday	Level 9 (incremental dump).
Tuesday	Level 8 (incremental dump).
Wednesday	Level 7 (incremental dump).
Thursday	Level 6 (incremental dump).
Friday	Level 5 (incremental dump).
Saturday	Level 4 (incremental dump).

Note that after the full backup on Sunday, a level 9 incremental dump is done the next day, and a successively lower dump level is done each day after that. This results in all the files that have changed since Sunday being backed up on every single incremental backup. Each incremental backup is thus larger than the previous; the backup contains all of the files from the previous incremental backup plus any files that have changed since then. This may seem wasteful of storage space on the backup tape, but it will save a lot of time and effort should there be a need to restore the file system.

For example, imagine that your hard drive crashed on Friday. After replacing the hard drive, you can restore the entire file system in two steps: restore the full backup from the prior Sunday and then the most recent incremental backup from Thursday. You can do this because Thursday's backup contains all of the files from Monday, Tuesday, and Wednesday's tape as well as the files that changed after

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

that. If the dump levels had progressed in positive order (level 1 for Monday, level 2 for Tuesday, and so on), *all* of the incremental backups would have to be restored in order to restore the file system to its most current state.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Automating Backups with cron

You can automate most of your backups with shell scripts and the cron daemon. Use the `su` command to become root, and then `cd` to the `/usr/local/bin` directory. Use any text editor to create a shell script called `backups.sh` that looks similar to the following:

```
#!/bin/sh
#
# backups.sh - A simple backup script, by Thad Phetteplace
#
# This script takes one parameter, the dump level.
# If the dump level is not provided, it is
# automatically set to zero. For level zero (full)
# dumps, rewind and eject the tape when done.
#

if [ $1 ]; then
        level=$1
else
        #
        # No dump level was provided, so set it
        # to zero
        #
        level="0"
fi

/sbin/dump $level'uf' /dev/nrft0 /
/sbin/dump $level'uf' /dev/nrft0 /home
/sbin/dump $level'uf' /dev/nrft0 /var
/sbin/dump $level'uf' /dev/nrft0 /usr

#
# If we are doing a full dump, rewind and eject
# the tape when done.
#
if [ $level = "0" ]; then
        #
        # Note: We should replace this with the /bin/st command
        # instead if we ever switch to a SCSI tape drive.
        #
        /bin/mt -f /dev/nrft0 rewind
        /bin/mt -f /dev/nrft0 offline
fi
```

You may choose to change the partitions being backed up to match your situation, but this script should otherwise work quite well for you. After saving and exiting the editor, change the permissions on

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

the file so that it is executable only by root:

chmod 700 backups.sh

You can now back up your entire system by running the `backups.sh` script when logged in as root. The script accepts the dump level as its only parameter. If you leave the parameter off, it will automatically assume a level zero dump. Thus, the following two commands are equivalent:

backups.sh

backups.sh 0

You may need to customize this script for your situation. For example, I am using the tape device `/dev/nrft0`. You may be using a different tape device. Whatever device you use, you should probably use the version of its device name that begins with a the letter *n*. That tells the system that after it finishes copying data to the tape, it should *not* rewind the tape. As an example, I used `/dev/nrft0` instead of `/dev/rst0` in the preceding script. If I had used `/dev/rst0`, each successive incremental backup would have overwritten the previous one.

Other things that you may change in this script include the partitions being backed up and the dump level at which the tape is ejected. It is not uncommon to eject the tape after the last incremental backup just before performing a full backup.

The most useful thing about this script is that you can easily configure your system to run it automatically. Simply add a few lines to the root `crontab` file, and the cron daemon will invoke the script on the days and times specified. While logged in as root, enter the `crontab` command with the `-e` option:

crontab -e

This opens the root `crontab` file in an editor. Add the following lines at the end of the file:

```
0 22 * * 0 /usr/local/bin/backup.sh 0
0 22 * * 1 /usr/local/bin/backup.sh 9
0 22 * * 2 /usr/local/bin/backup.sh 8
0 22 * * 3 /usr/local/bin/backup.sh 7
0 22 * * 4 /usr/local/bin/backup.sh 6
0 22 * * 4 /usr/local/bin/backup.sh 5
0 22 * * 5 /usr/local/bin/backup.sh 4
```

Save and exit the file. The cron daemon will now run the backup script at 10:00 p.m. (22:00 in military time) every day of the week. This example implements the dump schedule outlined earlier. A full dump is performed on Sunday, and the tape is ejected when it is done. A new tape should be loaded on Monday, and then incremental backups will be written to that same tape for the rest of the week. The next full dump will be written to the end of that tape, unless someone is around on Sunday to eject and replace the tape before 10:00 p.m.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Restoring Backed Up Files

The `restore` command is used to retrieve files from a backup tape or other medium that was created by `dump`. You can use `restore` to recover an entire file system or to interactively select individual files. It recovers files from the specified media and copies them into the current directory (the one you ran the `recover` command in), re-creating subdirectories as needed. Much as with the `dump` command, the first parameter passed to the `restore` command is a list of single character option codes, as shown in Table 13-5.

Table 13-5: Restore Command Options

Restore Options	Description
`-r`	Restore the entire dump archive.
`-C`	Compare the contents of the dump file with the files on the disk. This is used to check the success of a restore.
`-R`	Start the restore from a particular tape of a multitape backup. This is useful for restarting an interrupted restore.
`-X filelist`	Extract only specific files or directories from the archive. This option takes one argument, a list of files or directories to extract.
`-T file`	List the contents of the dump archive. If a file or directory is given as an argument, list only the occurrence of that file, directory, or anything within the directory.
`-i`	Restore files in interactive mode.
`-b blocksize`	Specify the block size of the dump in kilobytes. This option takes a numeric argument.
`-D filesystem`	Specify the name of the file system to be compared when using the `-C` option. The file system name is passed as an argument.
`-F script`	Specify the name of the dump archive to restore from. This option takes an alphanumeric argument.
`-h`	If this option is specified, `restore` re-creates directories marked for extraction but will not extract their contents.
`-m`	Files are extracted by inode number instead of name. This is generally not very useful.
`-N`	Instead of extracting files, print their names.
`-s file#`	Specify the dump file to start with on a multiple file tape. This takes a numeric argument.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 13-5: Restore Command Options

Restore Options	Description
`-T directory`	Tells `restore` where to write any temporary files. This is useful if you booted from a floppy (which has no space for temporary files).
`-v`	Run in verbose mode. This causes `restore` to print information about each file as it restores it.
`-y`	The `restore` command will always continue when it encounters a bad block, rather than asking you if you want to continue.

Restoring an entire file system

Let's return to our earlier example of the Friday disk crash. You installed a shiny new hard drive and your backup tapes are in hand. It is time to restore the files. For the purpose of this example, I assume that the crashed drive contained only the `/home` partition and that the Red Hat Linux operating system is still intact. If the crashed drive had contained the Red Hat Linux operating system, you would first have to reinstall Red Hat Linux before restoring the backup.

Before any files can be recovered to your new hard drive, an empty file system must be created on it. You will use the `mkfs` command to do this. The `mkfs` command can accept a variety of parameters, but usually you only need to supply the name of the device to create the file system on. Thus, to prepare the new hard drive type:

mkfs /dev/hda6

Alternatively, because our `/home` drive is listed in the `/etc/fstab` file, we can simply specify the `/home` mountpoint and `mkfs` will figure out the correct device. Thus, the preceding command could be replaced with this:

mkfs /home

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution You should, of course, exercise extreme caution when using the mkfs command. If you specify the wrong device name, you could unintentionally wipe out all data on an existing file sy

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

After creating a file system on your new disk, mount the partition to a temporary mount point.

mkdir /mnt/test

mount /dev/hda6 /mnt/test

This connects the new file system to the `/mnt/test` directory. Now change into the directory (`cd /mnt/test`) and use the `restore` command to recover the entire file system off of your backup tape. Of course, it is assumed that you have loaded the tape into the tape drive.

cd /mnt/test

restore rf /dev/nrft0

When the restore is finished, you can unmount the partition and remount it to the appropriate mount point. If you have restored the file system to a different physical partition than it was originally on, be sure to modify the `/etc/fstab` file appropriately so that the correct partition is mounted next time the system is rebooted.

Recovering individual files

The `restore` command can also be used to recover individual files and directories. By using `restore` in interactive mode, you can type a series of restore commands to selectively restore files. To run `restore` in interactive mode, use the `i` parameter instead of `r`:

restore if /dev/nrst0

The `restore` command will then read the file index from the backup tape and present you with a restore prompt. At this prompt, you can type the commands that enable you to select which directories and files to recover. You can navigate the directory structure of the backup index much the same way that you navigate an actual file system using a shell prompt. The interactive `restore` command even has its own version of the familiar `cd` and `ls` commands, as shown in Table 13-6.

Table 13-6: Interactive restore Commands

Command	Description
`add`	Add a file or directory to the list of files to be extracted. If a directory is marked for extraction, all of the directories and files within it will also be extracted.
`cd`	Change the current directory being viewed within the dump archive. Works similar to the `cd` command used at a shell prompt.
`delete`	Delete a file or directory from the list of files to be extracted. Deleting a directory from the list results in all of the files and directories within it also being deleted.
`extract`	Extract all of the marked files and directories from the archive and write them back to the file system.
`help`	List the available commands.
`ls`	List the contents of the current directory. If a directory name is provided as an argument, list the contents of that directory. Files or directories marked for extraction have a `*` character in front of them.
`pwd`	Print the full path name of the current directory of the dump archive.
`quit`	Exit the interactive restore program.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 13-6: Interactive restore Commands

Command	Description
`setmodes`	Do not restore the files; instead, set the modes of already existing files on the target disk to match the modes recorded in the dump file. This is useful for recovering from a restore that was prematurely aborted.
`verbose`	Toggles verbose output versus quiet output during the restore process. Verbose output mode will echo information to the screen for every file that is restored.

As an example, pretend that the user `joe` has accidentally deleted his Mail subdirectory from his home directory. Joe happens to be your boss, so it is urgent that you recover his files. Here is how you may go about it.

Load the appropriate tape into the tape drive and log in as root. Use the `cd` command to go to the top of the `/home` partition, and then run the `restore` program in interactive mode:

cd /home

restore if /dev/nrft0

Verify that you have the backup tape for the `/home` partition by entering the `ls` command. You should see something like the following list of directories, representing users that have home directories in `/home`:

restore > **ls**

.:

bob/ jane/ joe/ lost+found/ mary/ thad/

Yes, this is the home partition. Now change the current directory to Joe's home directory using the `cd` command. Type **ls** again to view the contents of his home directory.

restore > **cd joe**

restore > **ls**

./joe:

.mozilla/ Desktop/ report.html

.tcshrc Mail/ letter.txt

.xinitrc News/ www/

Now mark the Mail directory for extraction using the `add` command:

restore > **add Mail**

If you use the `ls` command again, you see that the Mail directory is preceded with an asterisk (`*`) character, which means it has been marked for extraction.

restore > **ls**

./joe:

.mozilla/ Desktop/ report.html

.tcshrc *Mail/ letter.txt

.xinitrc News/ www/

Now use the `extract` command to begin the extraction process. `restore` will prompt you for the number of the tape to start with. This is a single tape backup, so just enter the number `1`. When it prompts you to "`set owner/mode for '.'?`," answer yes by typing `y` and pressing Enter. `restore` will then set restore the file permissions (if necessary) of the directory it is restoring to. This isn't critical when extracting individual files like this, but you should always answer yes to this prompt when doing a full restore. Anyway, your screen should now contain the following:

restore > **extract**

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You have not read any tapes yet.

Unless you know which volume your file(s) are on you

should start with the last volume and work toward the first.

Specify next volume #: **1**

set owner/mode for '.'? [yn] **y**

restore >

At this point, the files have been recovered and you can exit the `restore` program by issuing the `quit` command. That's all there is to it. You now know the basics of using the `dump` and `restore` commands.

Configuring Amanda for network backups

Using Amanda (the Advanced Maryland Automatic Network Disk Archiver), you can use a single large-capacity tape drive on a server to back up files from multiple computers over a network. The Amanda package includes a variety of commands. The online man page for Amanda describes the commands as shown in Table 13-7.

Table 13-7: Backup Commands Used with Amanda

Command	Description
`amdump`	Do automatic Amanda backups. It is normally run by cron on a computer called the tape server host and requests backups of file systems located on backup clients. `amdump` backs up all disks in the disklist file to tape or, if there is a problem, to a special holding disk. After all backups are done, `amdump` sends mail reporting failures and successes.
`amflush`	Flush backups from the holding disk to tape. `amflush` is used after `amdump` has reported it could not write backups to tape for some reason. When this happens, backups stay in the holding disk. After the tape problem is corrected, run `amflush` to write backups from the holding disk to the tape.
`amcleanup`	Clean up after an interrupted `amdump`. This command is only needed if `amdump` was unable to complete for some reason, usually because the tape server host crashed while `amdump` was running.
`amrecover`	Provide an interactive interface to browse the Amanda index files and select which tapes to recover files from. `amrecover` can also run `amrestore` and the system restore program (such as tar) sometimes.
`amrestore`	Read an Amanda tape, searching for requested backups. `amrestore` is suitable for everything from interactive restores of single files to a full restore of all partitions on a failed disk.
`amlabel`	Write an Amanda format label onto a tape. All Amanda tapes must be labeled with `amlabel`. `amdump` and `amflush` will not write to an unlabeled tape.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 13-7: Backup Commands Used with Amanda

Command	Description
`amcheck`	Verify the correct tape is in the tape drive and that all file systems on all backup client systems are ready to be backed up. Can optionally be run by cron before `amdump`, so someone will get mail warning that backups will fail unless corrective action is taken.
`amadmin`	Take care of administrative tasks, such as finding out which tapes are needed to restore a file system, forcing hosts to do full backups of selected disks, and looking at schedule balance information.
`amtape`	Take care of tape changer control operations, such as loading particular tapes, ejecting tapes, and scanning the tape rack.
`amverify`	Check Amanda backup tapes for errors (GNU tar format backups only).
`amrmtape`	Delete a tape from the tape list and from the Amanda database.
`amstatus`	Give the status of a running `amdump`.

The `amdump` command is the one that you will use the most, but before you can get started, you need to configure a few things on both the backup server (the system with the tape drive) and the backup clients (the systems being backed up).

Creating Amanda directories

You need to create some directories to hold the Amanda configuration files and to provide a location to write Amanda log files. The configuration files go in an `/etc/amanda` directory, and the log files go in `/var/lib/amanda`. In both cases, you should log in as the amanda user and create subdirectories within those directories, one subdirectory for each backup schedule that you intend to run and an index file, as shown here:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note For security reasons, you should do all aman da administration as the aman da user. To do this, the root user should create a password for aman da by typing passwd aman da and entering

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

su - amanda

$ **mkdir -p /var/lib/amanda/normal**

$ **mkdir –p /var/lib/amanda/normal/index**

$ **mkdir -p /etc/amanda/normal**

For the purpose of this example, I have created only a `normal` backup configuration that backs up the data drives on several machines. You may also decide to create an upgrade backup configuration that backs up the operating system partitions. You could then run that backup before you perform any operating system upgrades.

You also need to specify a holding disk that Amanda can use to temporarily spool backups before it writes them to disk. This directory should have a lot of free space. I have a large `/home` partition on my server, so I created an Amanda directory there to use as a holding disk:

mkdir /home/amanda

chmod 700 /home/amanda

chown amanda /home/amanda

chgrp disk /home/amanda

Creating the amanda.conf file

Next, as the amanda user, you must create two configuration files for Amanda and store them in the `/etc/amanda/normal` directory: `amanda.conf` and `disklist`. You can start by copying samples of these files from the `/etc/amanda/DailySet1` directory as follows:

$ **cd /etc/amanda/DailySet1**

$ **cp amanda.conf disklist /etc/amanda/normal**

The `amanda.conf` file sets a variety of general configuration values, and the disklist file defines which machines and partitions to back up. The `amanda.conf` file can be rather complicated, but, fortunately, most of its values can be left at their default values. Here is a simplified `amanda.conf` file with some comments embedded in it to help explain things:

```
#
# amanda.conf - sample Amanda configuration file. This started life as
#               the actual config file in use at CS.UMD.EDU.


org "GLACI"       # your organization name for reports
mailto "amanda"   # space separated list of operators at your site
dumpuser "amanda"  # the user to run dumps under


# Specify tape device and/or tape changer. If you don't have a tape
# changer, and you don't want to use more than one tape per run of
# amdump, just comment out the definition of tpchanger.


runtapes 1             # number of tapes to be used in a single
                    # run of amdump
tapedev "/dev/nrft0"   # the no-rewind tape device to be used
rawtapedev "/dev/rft0"  # the raw device to be used (ftape only)


tapetype HP-DAT   # what kind of tape it is
       # (see tapetypes below)
labelstr "^normal[0-9][0-9]*$"   # label constraint all
                           # tapes must match
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
# Specify holding disks. These are used as a temporary
# staging area for dumps before they are written to tape and
# are recommended for most sites.

holdingdisk hd1 {
    comment "main holding disk"
    directory "/home/amanda"   # where the holding disk is
    use 290 Mb              # how much space can we use on it
    chunksize -1            # size of chunk
    }

# Note that, although the keyword below is infofile, it is
# only so for historic reasons, since now it is supposed to
# be a directory (unless you have selected some database
# format other than the `text' default)

infofile "/usr/adm/amanda/normal/curinfo"   # database DIRECTORY
logdir   "/usr/sdm/amanda/normal"      # log directory
indexdir "/usr/adm/amanda/normal/index"   # index directory

# tapetypes

# Define the type of tape you use here, and use it in "tapetype"
# above. Some typical types of tapes are included here. The
# tapetype tells amanda how many MB will fit on the tape, how
# big the filemarks are, and how fast the tape device is.

define tapetype HP-DAT {
    comment "DAT tape drives"
    # data provided by Rob Browning <rlb@cs.utexas.edu>
    length 1930 mbytes
    filemark 111 kbytes
    speed 468 kbytes
}

# dumptypes
#
# These are referred to by the disklist file.

define dumptype global {
    comment "Global definitions"
    # This is quite useful for setting global parameters, so you
    # don't have to type them everywhere.
}
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
define dumptype always-full {
    global
    comment "Full dump of this filesystem always"
    compress none
    priority high
    dumpcycle 0
}
```

This example `amanda.conf` was trimmed down from a larger example I copied from the `/etc/amanda/DailySet1` directory. The example `amanda.conf` file provides additional information on the available configuration options. Also, the online man page for Amanda should be helpful (type `man amanda` to read it). You can find more instructions in the `/usr/share/doc/amanda-server*` directory. Generally, you have to do the following:

Modify the org name for reports.
Change the device names set for tapedev and rawtapedev to match your tape device.
Select a tape type entry that is appropriate for your tape drive.
Change the name of the directory specified in the holding disk section to match the directory you created earlier.

Creating a disklist file

You also must create a `disklist` file in the `/etc/amanda/normal` directory. This simply contains a list of the systems and disk partitions to back up. The qualifier `always-full` is included on each entry to tell Amanda what type of backup to perform. It means to use full, rather than incremental, backups.

```
# sample Amanda2 disklist file
#
# File format is:
#
#       hostname diskdev dumptype [spindle [interface]]
#
# where the dumptypes are defined by you in amanda.conf.


dexter hda5 always-full
dexter hda6 always-full
dexter hda7 always-full
dexter hda8 always-full


daffy hda5 always-full
daffy hda6 always-full
daffy hda7 always-full
daffy hdb1 always-full
daffy hdb2 always-full
```

This example file backs up two systems, `dexter` and `daffy`. The order of the systems and the partitions is selected so that the most important data is backed up first. This way, if a tape drive becomes full, you have still managed to back up the most important data.

Adding Amanda network services

Amanda is designed to perform backups over a network. The following `amanda` services are defined in the `/etc/services` file:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

amanda 10080/udp

amanda 10080/tcp

amandaidx 10082/tcp

amidxtape 10083/tcp

On the amanda server

To offer these services to the network in Red Hat Linux, you need to configure the `xinetd` daemon to listen for those services. You do this by enabling the `amandaidx` and `amidxtape` services by typing the following (as root user):

chkconfig amidxtape on

chkconfig amandaidx on

This enables Amanda to accept requests from the client system and to start the backup process without any user intervention. You need to tell the `xinetd` daemon to reload the `/etc/xinetd.d` files before this change takes place. You can do this by typing the following as root user:

/etc/init.d/xinetd restart

On each amanda client

You next need to configure the `.amandahosts` file in the `/var/lib/amanda` directory on each computer (client) that the amanda server will backup from. This file should contain the fully qualified host and domain name of any backup servers that will connect to this client. When you begin, only your localhost is defined in this file as your backup server. To add another computer as a backup server, you could type the following (replacing *amandahost* with the name of the backup server, while you are logged in as the `amanda` user):

$ **echo *amandahost* >> /var/lib/amanda/.amandahosts**

You also need to make sure that the amanda client daemon is configured to run on the client. You do this by enabling the `amanda` service by typing the following (as root user) :

chkconfig amanda on

This enables the amanda client to communicate with the amanda server. You need to tell the `xinetd` daemon to reload the `/etc/xinetd.d` files before this change takes place. You can do this by typing the following as root user:

/etc/init.d/xinetd restart

Performing an Amanda backup

Now that everything is configured, you are ready to perform an Amanda backup. While logged in as root, type the following command:

/usr/sbin/amdump normal

This runs the `amdump` command and tells it to read the configuration files it finds in the `/etc/amanda/normal` directory created earlier. It then works its way down the list of systems and partitions in the `disklist` file, backing up each partition in the order it occurs. The results of the `amdump` are written to the `/var/lib/amanda/normal` directory. Read the files you find there to check on the results of the backup. (See the previous section on how to create a `disklist` file to understand the process that `amdump` goes through.)

You can, of course, automate this process with `cron`. To create an `amdump` schedule similar to the regular dump schedule discussed in an earlier section, do the following. While logged in as root, enter the `crontab` command with the `-e` option:

crontab -e

This opens the root `crontab` file in an editor. Add the following lines to the end of the file:

0 22 * * 0 /usr/sbin/amdump normal

0 22 * * 1 /usr/sbin/amdump incremental

0 22 * * 2 /usr/sbin/amdump incremental

0 22 * * 3 /usr/sbin/amdump incremental

0 22 * * 4 /usr/sbin/amdump incremental

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
0 22 * * 5 /usr/sbin/amdump incremental
0 22 * * 6 /usr/sbin/amdump incremental
```

Save and exit the file. The cron daemon will now run `amdump` at 10:00 p.m. (22:00 in military time) every day of the week. This example assumes that a second incremental configuration has been created. You can do this by creating a subdirectory named `incremental` under `/etc/amanda` and populating it with appropriately modified `amanda.conf` and `disklist` files. You must also create a subdirectory named `incremental` under `/usr/adm/amanda` so that `amanda` has somewhere to write the logfiles for this configuration.

It may be a bit of work to get it all in place, but once you do, Amanda can make your network backups much easier to manage. It may be overkill for a small office, but in a large enterprise network situation, it enables Red Hat Linux to act as a powerful backup server.

Team LiB

PREVIOUS NEXT

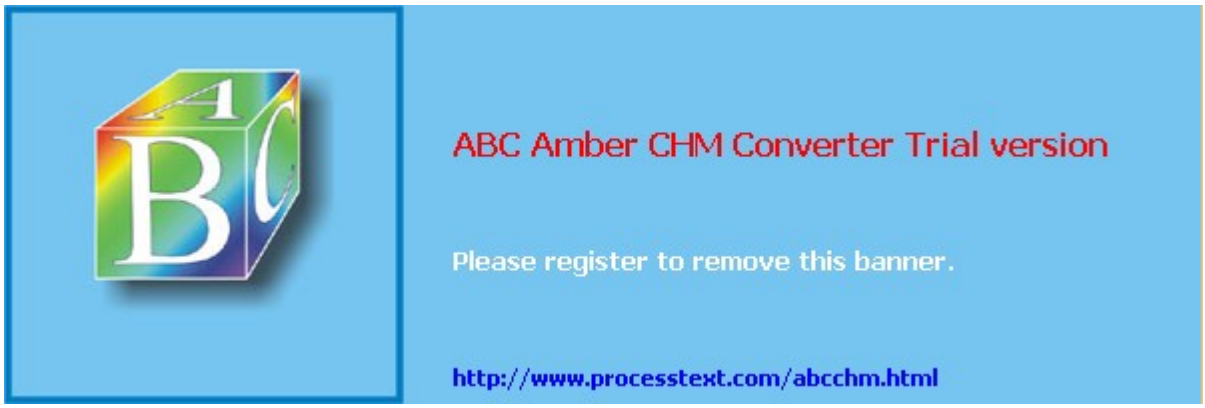

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using the pax Archiving Tool

Over the years, a variety of UNIX operating systems have arisen, resulting in a variety of similar but incompatible file archiving formats. Even tools that go by the same name may use slightly different storage formats on different systems. This can lead to big problems when trying to archive and retrieve data in a multiplatform environment. Fortunately, there is a solution.

The `pax` program is a POSIX standard utility that can read and write a wide variety of archive formats. An RPM package for `pax` is included with Red Hat Linux. If it is not already installed, copy the `pax-*` RPM file from your distribution media (CD #1), or download it from a Red Hat Linux FTP site, and then use the `rpm` command to install it.

rpm -i pax-*

Remember you need to be logged in as root when installing software with the `rpm` command.

Pax takes a variety of command-line options. The last parameter is usually the file or directory to archive. You may use wildcard characters such as "*" or "?" to specify multiple files or directories. The options you will use most often include the `-r` and `-w` parameters for specifying when you are reading or writing an archive. These are usually used in conjunction with the `-f` parameter, which is used to specify the name of the archive file.

By using `pax` parameters in different combinations, it is possible to extract an archive, create an archive, list the contents of an archive, or even copy an entire directory hierarchy from one location to another. Table 13-8 shows a few examples of the `pax` command in action.

Table 13-8: Examples of pax Use

Pax Command	Description
`pax -f myfiles`	List the contents of the archive named `myfiles`.
`pax -r -f myfiles`	Extract the contents of the archive named `myfiles`.
`pax -w -f myfiles /etc`	Create an archive named `myfiles` containing everything within the `/etc` directory.
`pax -w -f myfiles *.txt`	Archive all of the files in the current directory that have a `.txt` file extension.
`pax -r -w /olddir /newdir`	Copy the entire contents of the directory `/oldir` into a new directory called `/newdir`.
`pax -w -B 1440000 -f /dev/fd0 *`	Archive the contents of the current directory onto multiple floppy disks.
`pax -w -x cpio -f myfiles *`	Archive the contents of the current directory into an archive file named `myfiles` using the `cpio` format.
`pax -r -U mary -f backups`	Extract all of the files owned by user `mary` from the archive named `backups`.

Note that by leaving off both the `-r` and `-w` options, you cause `pax` to simply list the contents of the archive. If you specify both the `-r` and `-w` options, then you should leave off the `-f` option and supply source and destination directories instead. This will cause the source directory to be completely cloned in the specified destination directory.

You can use additional parameters to further modify `pax`'s behavior. For example, you may use the `-x` option in conjunction with the `-w` option to specify the specific archive type to create, or you may use the `-B` option to specify the number of bytes to write to each volume of a multi-volume archive. Table 13-9 briefly describes the many optional parameters to the `pax` command.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 13-9: Options to pax

Pax Options	Description
`-r`	Read files from an archive.
`-w`	Write files to an archive.
`-a`	Append files to a previously created archive.
`-b` *`blocksize`*	Specify the archive's data block size. It must be a multiple of 512.
`-c`	Match all files except those that match the specified pattern.
`-d`	Match filename wildcards against file or directory names only, not the complete path.
`-f` *`archive`*	Specify the name of the archive.
`-i`	Interactively rename files when archiving.
`-k`	Do not overwrite existing files.
`-l`	Link files with hard links when in copy mode (`-r -w`).
`-n`	Match only the first file that matches the supplied pattern.
`-o` *`options`*	Extra options specific to the archiving format used.
`-p` *`string`*	Specify the file characteristics to retain when archiving or copying. Read the `pax` man page for more information on this option.
`-s` *`replstr`*	Modify the archived filenames using the supplied regular expression.
`-t`	Preserve the access times of archived files.
`-u`	Do not overwrite files with older versions.
`-v`	Provide verbose output when running.
`-x` *`format`*	Specify format of the archive. Valid formats include `cpio`, `bcpio`, `sv4cpio`, `sv4crc`, `tar`, and `ustar`. The default is to use `ustar` when creating an archive. Pax will automatically determine the correct file type when reading an archive.
`-z`	Indicates that `gzip` should be used to compress/decompress the archive.
`-B` *`bytes`*	Specify the number of bytes per archive volume. Use this option to create multivolume archives on removable media.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 13-9: Options to pax

Pax Options	Description
`-D`	Do not overwrite existing files with files that have an older inode modification time.
`-E` *limit*	Limit the number of times `pax` will retry on encountering a read or write error.
`-G` *group*	Select files based on a group name or GID. To select by GID, place a # sign in front of the group number.
`-H`	Follow only command-line symbolic links while performing a physical file system traversal.
`-L`	Follow all symbolic links when traversing a directory hierarchy.
`-P`	Do not follow symbolic links. This is the default.
`-T` *time*	Select files based on their modification time. Read the `pax` man page for complete discussion of this parameter's syntax.
`-U` *user*	Select files based on the owner's user name, or by UID with a # sign in front of it.
`-X`	Do not traverse into directories that reside on a different device.
`-Y`	This option is the same as the `-D` option, except that the inode change time is checked using the pathname created after all the filename modifications have completed.
`-Z`	This option is the same as the `-u` option, except that the modification time is checked using the pathname created after all the filename modifications have completed.

As you can see, `pax` is a very flexible and powerful archiving tool. It can be particularly helpful in migrating data from older legacy systems to your new Linux system. When you are faced with the task of recovering archived data from an antiquated or even nonfunctioning UNIX system, the multiple file format support of `pax` can be a literal lifesaver.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Summary

I hope that you never experience a major hard drive crash, but if you ever do, the effort of making backups will repay itself many times over. A variety of low-cost backup hardware is available to use with your Red Hat Linux system. The traditional tape drive is an excellent choice for backing up large amounts of data. If long-term archiving of data is needed, a writable DVD or CD drive is a good choice. If minimizing downtime is your main concern, mirroring data to a second hard drive is another smart choice. As for the tools you choose to do your backups, the `dump` and `restore` commands are old favorites, dating back from the early UNIX days. The Amanda backup facility is excellent for network backups. If you are dealing with backups in several different formats, the `pax` command might be helpful.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Hacker versus Cracker

In short, a hacker is someone who programs creatively and usually for the pure enjoyment of it (most programmers who work on Linux are hackers in this sense). The correct term for someone who breaks into computer systems is a cracker.

There are many types of crackers, ranging from professional computer criminals to the hobbyist types that break into computers for the thrill. The growth of the cracker problem has kept pace with the growth of the Internet. A new, younger generation of cracker is emerging. These teenage pseudo-crackers do not have all the knowledge and skill of their true cracker counterparts, but they have access to a growing number of cracker tools that automate the breaking of a system's security.

By using programs and scripts created by more advanced crackers, youngsters can often break into systems without really knowing the details of how it is done. Because they are usually young and mostly dependent on tools provided by others, they are sometimes referred to as "scriptkiddies." Make no mistake, if your system is not properly secured, scriptkiddies can do just as much damage as any other cracker.

Whatever you call them, crackers pose a serious risk to anyone connecting a computer to the Internet. Their reasons for breaking into systems are varied; some hope to steal financial information, others wish to gain bragging rights among their peers.

Often, a system is broken into solely for use as a jumping-off point to launch further attacks on other systems. In some cases, the damage may be as little as an altered Web page, the Internet equivalent of graffiti. In other cases, the cracker may wipe out your entire hard drive to cover his or her tracks. Fortunately, there are ways to protect yourself.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Understanding Attack Techniques

Attacks on computing systems take on different forms, depending on the goal and resources of the attacker. Some attackers desire to be disruptive, while others desire to infiltrate your machines and utilize the resources for their own nefarious purposes. Still others are targeting your data for financial gain or blackmail. Here are three major categories of attacks:

Denial of Service (DOS) — The easiest attacks to perpetrate are Denial of Service attacks. The primary purpose of these attacks is to disrupt the activities of a remote site by overloading it with irrelevant data. DOS attacks can be as simple as sending thousands of page requests per second at a Web site. These types of attacks are easy to perpetrate and easy to protect against. Once you have a handle on where the attack is coming from, a simple phone call to the perpetrator's ISP will get the problem solved.

Distributed Denial of Service (DDOS) — More advanced DOS attacks are called Distributed Denial of Service attacks. DDOS attacks are much harder to perpetrate and nearly impossible to stop. In this form of attack, an attacker takes control of hundreds or even thousands of weakly secured Internet connected computers. The attacker then directs them in unison to send a stream of irrelevant data to a single Internet host. The result is that the power of one attacker is magnified thousands of times. Instead of an attack coming from one direction, as is the case in a normal DOS, it comes from thousands of directions at once. The best defense against DDOS attack is to contact your ISP to see if it can filter traffic at its border routers.

Many people use the excuse, "I have nothing on my machine anyone would want" to avoid having to consider security. The problem with this argument is that attackers have a lot of reasons to use your machine. The attacker can turn your machine into an agent for later use in a DDOS attack. More than once, authorities have shown up at the door of a dumbfounded computer user asking questions about threats originating from their computer. By ignoring security, the owners have opened themselves up to a great deal of liability.

Intrusion attacks — To remotely use the resources of a target machine, attackers must first look for an opening to exploit. In the absence of inside information such as passwords or encryption keys, they must scan the target machine to see what services are offered. Perhaps one of the services is weakly secured and the attacker can use some known exploit to finagle his way in.

A tool called `nmap` is generally considered the best way to scan a host for services (note that `nmap` is a tool that can be used for good and bad). Once the attacker has a list of the available services running on his target, he needs to find a way to trick one of those services into letting him have privileged access to the system. Usually, this is done with a program called an *exploit*.

While DOS attacks are disruptive, intrusion type attacks are the most damaging. The reasons are varied, but the result is always the same. An uninvited guest is now taking up residence on your machine and is using it in a way you have no control over.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Protecting Against Denial-of-Service Attacks

As explained earlier, a denial-of-service attack attempts to crash your computer or at least degrade its performance to an unusable level. There are a variety of denial-of-service exploits. Most try to overload some system resource, such as your available disk space or your Internet connection. Some common attacks and defenses are discussed in the following sections.

Mailbombing

Mailbombing is the practice of sending so much e-mail to a particular user or system that the computer's hard drive becomes full. There are several ways to protect yourself from mailbombing. You can use the Procmail e-mail-filtering tool or configure your sendmail daemon.

Cross-Reference See Chapter 19 for a more complete description of sendmail.

Blocking mail with Procmail

The Procmail e-mail-filtering tool is installed by default with Red Hat Linux and is tightly integrated with the sendmail e-mail daemon; thus, it can be used to selectively block or filter out specific types of e-mail. You can learn more about Procmail at the Procmail Web site `www.procmail.org`.

To enable Procmail for your user account, create a `.procmailrc` file in your home directory. The file should be mode 0600 (readable by you but nobody else). Type the following, replacing *evilmailer* with the actual e-mail address that is mailbombing you.

```
# Delete mail from evilmailer
:0
* ^From.*evilmailer
/dev/null
```

The Procmail recipe looks for the `From` line at the start of each e-mail to see if it includes the string `evilmailer`. If it does, the message is sent to `/dev/null` (effectively throwing it away).

Blocking mail with sendmail

The Procmail e-mail tool works quite well when only one user is being mailbombed. If, however, the mailbombing affects many users, you should probably configure your sendmail daemon to block all e-mail from the mailbomber. Do this by adding the mailbomber's e-mail address or system name to the `access` file located in the `/etc/mail` directory.

Each line of the `access` file contains an e-mail address, host name, domain, or IP address followed by a tab and then a keyword specifying what action to take when that entity sends you a message. Valid keywords are OK, RELAY, REJECT, DISCARD, and ERROR. Using the REJECT keyword will cause a sender's e-mail to be bounced back with an error message. The keyword DISCARD will cause the message to be silently dropped without sending an error back. You can even return a custom error message by using the ERROR keyword.

Thus, an example `/etc/mail/access file` may look similar to this:

```
# Check the /usr/share/doc/sendmail/README.cf file for a description
# of the format of this file. (search for access_db in that file)
# The /usr/share/doc/sendmail/README.cf is part of the sendmail-doc
# package.
#
# by default we allow relaying from localhost...
localhost.localdomain         RELAY
localhost                     RELAY
127.0.0.1                     RELAY
#
# Senders we want to Block
#
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
evilmailer@yahoo.com   REJECT
stimpy.glaci.com       REJECT
cyberpromo.com         DISCARD
199.170.176.99         ERROR:"550 Die Spammer Scum!"
199.170.177            ERROR:"550 Email Refused"
```

As with most Linux configuration files, lines that begin with a # pound sign are comments. Our list of blocked spammers is at the end of this example file. Note that the address to block can be a complete e-mail address, a full host name, a domain only, an IP address, or a subnet.
To block a particular e-mail address or host from mailbombing you, log in to your system as root, edit the `/etc/mail/access` file, and add a line to DISCARD mail from the offending sender.

After saving the file and exiting the editor, you must convert the access file into a hash indexed database called access.db. The database is updated automatically the next time sendmail starts. Or you can convert the database immediately, as follows:

```
# cd /etc/mail
# make
```

Sendmail should now discard e-mail from the addresses you added.

Spam relaying
Another way in which your e-mail services can be abused is by having your system used as a Spam Relay. *Spam* refers to the unsolicited junk e-mail that has become a common occurrence on the Internet. Spammers often deliver their annoying messages from a normal dial-up Internet account. They need some kind of high-capacity e-mail server to accept and buffer the payload of messages. They deliver the spam to the server all in one huge batch and then log off, letting the server do the work of delivering the messages to the many victims.

Naturally, no self-respecting Internet Service Provider will cooperate with this action, so spammers resort to hijacking servers at another ISP to do the dirty work. Having your mailserver hijacked to act as a spam relay can have a devastating effect on your system and your reputation. Fortunately, mail relaying is deactivated by default on Red Hat Linux installations. This is one security issue that you will not have to attend to.
You can allow specific hosts or domains to relay mail through your system by adding those senders to your `/etc/mail/access` file with keyword RELAY. By default, relaying is only allowed from the local host. Refer to Chapter 19, as well as the sendmail documentation, for more information.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip One package you might consider using to filter out spam on your mail server is Spamassassin. Spamassassin examines the text of incoming mail messages and attempts to filter

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Smurf amplification attack

Smurfing refers to a particular type of denial-of-service attack aimed at flooding your Internet connection. It can be a difficult attack to defend against because it is not easy to trace the attack to the attacker. Here is how smurfing works.

The attack makes use of the ICMP protocol, a service intended for checking the speed and availability of network connections. Using the `ping` command, you can send a network packet from your computer to another computer on the Internet. The remote computer will recognize the packet as an ICMP request and echo a reply packet to your computer. Your computer can then print a message revealing that the remote system is up and telling you how long it took to reply to the ping.

A smurfing attack uses a malformed ICMP request to bury your computer in network traffic. The attacker does this by bouncing a ping request off an unwitting third party in such a way that the reply is duplicated dozens or even hundreds of times. An organization with a fast Internet connection and a large number of computers is used as the relay. The destination address of the ping is set to an entire subnet instead of a single host. The return address is forged to be your machine's address instead of the actual sender. When the ICMP packet arrives at the unwitting relay's network, every host on that subnet replies to the ping! Furthermore, they reply to your computer instead of to the actual sender. If the relay's network has hundreds of computers, your Internet connection can be quickly flooded.

The best fix is to contact the organization being used as a relay and inform them of the abuse. Usually, they need only to reconfigure their Internet router to stop any future attacks. If the organization is uncooperative, you can minimize the effect of the attack by blocking the ICMP protocol on your router. This will at least keep the traffic off your internal network. If you can convince your ISP to block ICMP packets aimed at your network, it will help even more.

Team LiB

PREVIOUS NEXT

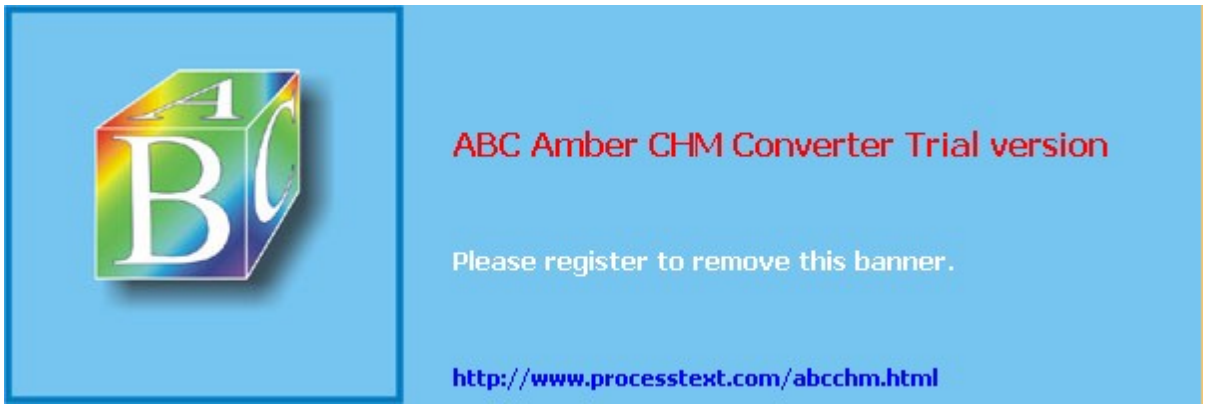

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Protecting Against Distributed DOS Attacks

DDOS attacks are much harder to initiate and nearly impossible to stop. A DDOS attack begins with the penetration of hundreds or even thousands of weakly secured machines. These machines can then be directed to attack a single host based on the whims of the attacker.
With the advent of DSL and cable modem, millions of people are enjoying Internet access with virtually no speed restrictions. In their rush to get online, many of those people neglect even the most basic security. Since the vast majority of these people run Microsoft operating systems, they tend to get hit with worms and viruses rather quickly. Once the machine has been infiltrated, quite often the worm or virus installs a program on the victim's machine that instructs it to quietly *call home* and announce that it is now ready to do *the master's bidding*.
At the whim of the master, the infected machines can now be used to focus a concentrated stream of garbage data at a selected host. In concert with thousands of other infected machines, a *scriptkiddie* now has the power take down nearly any site on the Internet.

Detecting a DDOS is similar to detecting a DOS attack. One or more of the following signs are likely to be present:

- Sustained saturated data link
- No reduction in link saturation during off-peak hours
- Hundreds or even thousands of simultaneous network connections
- Extremely slow system performance

To determine if your data link is saturated, the act of pinging an outside host can tell much of the story. Much higher than usual latency is a dead giveaway. Normal ping latency (that is, the time it takes for a ping response to come back from a remote host) looks like the following:

```
# ping www.example.com
PING www.example.com (192.0.34.166) from 10.0.0.11: 56(84) bytes of data
64 bytes from 192.0.34.166: icmp_seq=1 ttl=49 time=40.1 ms
64 bytes from 192.0.34.166: icmp_seq=2 ttl=49 time=42.5 ms
64 bytes from 192.0.34.166: icmp_seq=3 ttl=49 time=39.5 ms
64 bytes from 192.0.34.166: icmp_seq=4 ttl=49 time=38.4 ms
64 bytes from 192.0.34.166: icmp_seq=5 ttl=49 time=39.0 ms

--- www.example.com ping statistics ---
5 packets transmitted, 5 received, 0% loss, time 4035ms
rtt min/avg/max/mdev = 38.472/39.971/42.584/1.432 ms
```

In the preceding example, the average time for a ping packet to make the round trip was about 39 thousandths of a second.

A ping to a nearly saturated link will look like the following:

```
# ping www.example.com
PING www.example.com (192.0.34.166): from 10.0.0.11: 56(84)bytes of data
64 bytes from 192.0.34.166: icmp_seq=1 ttl=62 time=1252 ms
64 bytes from 192.0.34.166: icmp_seq=2 ttl=62 time=1218 ms
64 bytes from 192.0.34.166: icmp_seq=3 ttl=62 time=1290 ms
64 bytes from 192.0.34.166: icmp_seq=4 ttl=62 time=1288 ms
64 bytes from 192.0.34.166: icmp_seq=5 ttl=62 time=1241 ms

--- www.example.com ping statistics ---
6 packets transmitted, 5 received, 0% loss, time 5032ms
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
rtt min/avg/max/mdev = 1218.059/1258.384/1290.861/28.000 ms
```

In this example, a ping packet took, on average, 1.3 seconds to make the round trip. From the first example to the second example, latency increased by a factor of 31! A data link that goes from working normally to slowing down by a factor of 31 is a clear sign that link utilization should be investigated.
For a more accurate measure of data throughput, a tool such as `ttcp` can be used. To test your connection with `ttcp` you must have installed the `ttcp` RPM package on machines inside *and* outside of your network. (The ttcp package comes on CD #3 included with this book.) If you are not sure if the package is installed, simply type `ttcp` at a command prompt. You should see something like the following:

```
# ttcp
Usage: ttcp -t [-options] host [ < in ]
       ttcp -r [-options > out]
Common options:
    -l ##     length of bufs read from or written to network (default 8192)
    -u     use UDP instead of TCP
    -p ##     port number to send to or listen at (default 5001)
    -s     -t: source a pattern to network
           -r: sink (discard) all data from network
    -A     align the start of buffers to this modulus (default 16384)
    -O     start buffers at this offset from the modulus (default 0)
    -v     verbose: print more statistics
    -d     set SO_DEBUG socket option
    -b ##     set socket buffer size (if supported)
    -f X     format for rate: k,K = kilo{bit,byte}; m,M = mega; g,G = giga
Options specific to -t:
    -n##     number of source bufs written to network (default 2048)
    -D     don't buffer TCP writes (sets TCP_NODELAY socket option)
Options specific to -r:
    -B     for -s, only output full blocks as specified by -l (for TAR)
    -T     "touch": access each byte as it's read
```

The first step is to start up a receiver process on the server machine:

```
# ttcp -rs
ttcp-r: buflen=8192, nbuf=2048, align=16384/0, port=5001  tcp
ttcp-r: socket
```

The `-r` flag denotes that the server machine will be the receiver. The `-s` flag, in conjunction with the `-r` flag, tells `ttcp` that we want to ignore any received data.
The next step is to have someone outside of your data link, with a network link close to the same speed as yours, set up a `ttcp` sending process:

```
# ttcp -ts server.example.com
ttcp-t: buflen=8192, nbuf=2048, align=16384/0, port=5001  tcp  ->
server.example.com
ttcp-t: socket
ttcp-t: connect
```

Let the process run for a few minutes and then press Ctrl+C on the transmitting side to stop the testing. The receiving side will then take a moment to calculate and present the results:

```
# ttcp -rs
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ttcp-r: buflen=8192, nbuf=2048, align=16384/0, port=5001 tcp

ttcp-r: socket

ttcp-r: accept from 64.223.17.21

ttcp-r: 2102496 bytes in 70.02 real seconds = 29.32 KB/sec +++

ttcp-r: 1226 I/O calls, msec/call = 58.49, calls/sec = 17.51

ttcp-r: 0.0user 0.0sys 1:10real 0% 0i+0d 0maxrss 0+2pf 0+0csw

In this example, the average bandwidth between the two hosts was 29.32 kilobytes per second. On a link suffering from a DDOS, this number would be a mere fraction of the actual bandwidth the data link is rated for.
If the data link is indeed saturated, the next step is to determine where the connections are coming from. A very effective way of doing this is with the `netstat` command, which is included as part of the base Red Hat Linux install. Type the following to see connection information:

netstat –tupn
Table 14-1 describes each of the `netstat` parameters used here.

Table 14-1: netstat Parameters

Parameter	Description
`-t, --tcp`	Show TCP socket connections.
`-u, --udp`	Show UDP socket connections.
`-p, --program`	Show the PID and name of the program to which each socket belongs.
`-n, --numeric`	Show numerical address instead of trying to determine symbolic host, port, or user names.

The following is an example of what the output might look like:

Active Internet connections (w/o servers)

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State	PID/Program name
tcp	0	0	65.213.7.96:22	13.29.132.19:12545	ESTABLISHED	32376/sshd
tcp	0	224	65.213.7.96:22	13.29.210.13:29250	ESTABLISHED	13858/sshd
tcp	0	0	65.213.7.96:6667	13.29.194.190:33452	ESTABLISHED	1870/ircd
tcp	0	0	65.213.7.96:6667	216.39.144.152:42709	ESTABLISHED	1870/ircd
tcp	0	0	65.213.7.96:42352	67.113.1.99:53	TIME_WAIT	-
tcp	0	0	65.213.7.96:42354	83.152.6.9:113	TIME_WAIT	-
tcp	0	0	65.213.7.96:42351	83.152.6.9:113	TIME_WAIT	-
tcp	0	0	127.0.0.1:42355	127.0.0.1:783	TIME_WAIT	-
tcp	0	0	127.0.0.1:783	127.0.0.1:42353	TIME_WAIT	-
tcp	0	0	65.213.7.96:42348	19.15.11.1:25	TIME_WAIT	-

The output is organized into columns defined as follows:

Proto — Protocol used by the socket.
Recv-Q — The number of bytes not yet copied by the user program attached to this socket.
Send-Q — The number of bytes not acknowledged by the host.
Local Address — Address and port number of the local end of the socket.
Foreign Address — Address and port number of the remote end of the socket.
State — Current state of the socket. Table 14-2 provides a list of socket states.
PID/Program name — Process ID and program name of the process that owns the socket.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 14-2: Socket States

State	Description
ESTABLISHED	Socket has an established connection.
SYN_SENT	Socket actively trying to establish a connection.
SYN_RECV	Connection request received from the network.
FIN_WAIT1	Socket closed and shutting down.
FIN_WAIT2	Socket is waiting for remote end to shut down.
TIME_WAIT	Socket is waiting after closing to handle packets still in the network.
CLOSED	Socket is not being used.
CLOSE_WAIT	The remote end has shutdown, waiting for the socket to close.
LAST_ACK	The remote end has shut down, and the socket is closed, waiting for acknowledgement.
LISTEN	Socket is waiting for an incoming connection.
CLOSING	Both sides of the connection are shut down but not all of our data has been sent.
UNKNOWN	The state of the socket is unknown.

During a DOS attack, the foreign address is usually the same for each connection. In this case, it is a simple matter of typing the foreign IP address into the search form over at http://www.arin.net/whois/ so you can alert your ISP.

During a DDOS, the foreign address will likely be different for each connection. In this case, it is impossible to track down all of the offenders, as there will likely be thousands of them. The best way to defend yourself is to contact your Internet provider and see if it can filter the traffic at its border routers.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Protecting Against Intrusion Attacks

Crackers have a wide variety of tools and techniques to assist them in breaking into your computer. Intrusion attacks focus on exploiting weaknesses in your security, so the crackers can take more control of your system (and potentially do more damage) than they could from the outside.

Fortunately, there are many tools and techniques for combating intrusion attacks. This section discusses the most common break-in methods and the tools available to protect your system. Though the examples shown are specific to Red Hat Linux systems, the tools and techniques are generally applicable to any other Linux or UNIX-like operating system.

Cross-Reference

The tripwire package, which was recently dropped from the Red Hat Linux distribution, was a good tool for detecting whether intrusion attacks have taken place. We removed the description of tripwire from this edition. However, if you are interested in installing the package on your own, you can find the description of tripwire at the Wiley Web site: www.wiley.com/legacy/compbooks/negus)

Evaluating access to network services

Red Hat Linux and its UNIX kin provide many network services, and with them many avenues for cracker attacks. You should know these services and how to limit access to them.
So what do I mean by a network service? Basically, I am referring to any task that the computer performs that requires it to send and receive information over the network using some predefined set of rules. Routing e-mail is a network service. So is serving Web pages. Your Linux box has the potential to provide thousands of services. Many of them are listed in the `/etc/services` file. Look at a snippet of that file:

```
# /etc/services:
# service-name  port/protocol  [aliases ...]   [# comment]
chargen         19/tcp          ttytst source
chargen         19/udp          ttytst source
ftp-data        20/tcp
ftp-data        20/udp
# 21 is registered to ftp, but also used by fsp
ftp             21/tcp
ftp             21/udp          fsp fspd
ssh             22/tcp                  # SSH Remote Login Protocol
ssh             22/udp                  # SSH Remote Login Protocol
telnet          23/tcp
telnet          23/udp
# 24 - private mail system
smtp            25/tcp          mail
```

After comment lines, you will notice three columns of information. The left column contains the name of each service. The middle column defines the port number and protocol type used for that service. The rightmost field contains an optional alias or list of aliases for the service.

As an example, let us examine the last entry in the above file snippet. It describes the SMTP (Simple Mail Transfer Protocol) service, which is the service used for delivering e-mail over the Internet. The middle column contains the text 25/tcp, which tells us that the SMTP protocol uses port 25 and uses the Transmission Control Protocol (TCP) as its protocol type.

So, what exactly is a port number? It is a unique number that has been set aside for a particular network service. It allows network connections to be properly routed to the software that handles that service. For example, when an e-mail message is delivered from some other computer to your Linux

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

box, the remote system must first establish a network connection with your system. Your computer receives the connection request, examines it, sees it labeled for port 25, and thus knows that the connection should be handed to the program that handles e-mail (which happens to be sendmail).

Note A program that stays quietly in the background handling service requests (such as sendmail) is called a daemon. Usually, daemons are started automatically when your system boots up, and they keep running until your system is shut down. Daemons may also be started on an as-needed basis by xinetd, a special daemon that listens on a large number of port numbers, then launches the requested process.

I mentioned that SMTP uses the TCP protocol. Some services use UDP, the User Datagram Protocol. All you really need to know about TCP and UDP (for the purpose of this security discussion) is that they provide different ways of packaging the information sent over a network connection. A TCP connection provides error detection and retransmission of lost data. UDP doesn't check to ensure that the data arrived complete and intact; it is meant as a fast way to send non-critical information.

Disabling network services

Although there are hundreds of services (listed in `/etc/services`) that potentially could be available and subject to attack on your Red Hat Linux system, in reality only a few dozen services are installed and only a handful of those are on by default. Most network services are started by either the `xinetd` process or by a start-up script in the `/etc/init.d` directory.

xinetd is a daemon that listens on a great number of network port numbers. When a connection is made to a particular port number, xinetd automatically starts the appropriate program for that service and hands the connection to it.

The configuration file `/etc/xinetd.conf` is used to provide default settings for the `xinetd` server. The directory `/etc/xinetd.d` contains files telling `xinetd` what ports to listen on and what programs to start. Each file contains configuration information for a single service, and the file is usually named after the service it configures. For example, to enable the Post Office Protocol (POP) service, edit the `ipop3` file and look for a section similar to the following:

```
service pop3
{
    socket_type          = stream
    wait                 = no
    user                 = root
    server               = /usr/sbin/ipop3d
    log_on_success       += HOST DURATION
    log_on_failure       += HOST
    disable              = yes
}
```

Note that the first line of this example identifies the service as `pop3`. This exactly matches the service name listed in the `/etc/services` file. You can see that the service is off by default (`disable = yes`). To enable POP services, change the line to read `disable = no` instead. Thus, the preceding example with POP services enabled would look like this:

```
service pop3
{
    socket_type          = stream
    wait                 = no
    user                 = root
    server               = /usr/sbin/ipop3d
    log_on_success       += HOST DURATION
    log_on_failure       += HOST
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
    disable         = no
}
```

Because most services are disabled by default, your computer is only as insecure as you make it. You can double-check that insecure services, such as `rlogin` and `rsh`, are also disabled by making sure that `disabled = yes` is set in the `/etc/xinetd.d/rlogin` and `rsh` files.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip You can make the remote login service active but disable the use of the /etc/host.equiv and .rhosts files, requiring rlogin to always prompt for a password. Rather than disabling the ser

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You now need to send a signal to the `xinetd` process to tell it to reload its configuration file. The quickest way to do that is to restart the `xinetd` service. As the root user, type the following from a shell:

```
# service xinetd restart
Stopping xinetd:          [ OK ]
Starting xinetd:          [ OK ]
```

That's it — you have enabled the `ipop3` service. Provided that you have properly configured your mail server (see Chapter 19), clients should now be able to get their mail from your computer.

Using TCP wrappers

Completely disabling an unused service is fine, but what about the services that you really need? How can you selectively grant and deny access to these services? In previous versions of Red Hat Linux, the TCP wrapper daemon (`tcpd`) was used to facilitate this sort of selective access. In the current version of Red Hat Linux, TCP wrapper support has been integrated into `xinetd`. `xinetd` will look at the files `/etc/hosts.allow` and `/etc/hosts.deny` to determine when a particular connection should be granted or refused. It scans through the `hosts.allow` and `hosts.deny` files and stops when it finds an entry that matches the IP address of the connecting machine. These checks are made when connection attempts occur:

If the address is listed in the `hosts.allow` file, the connection is allowed and `hosts.deny` is not checked.
Otherwise, if the address is in `hosts.deny`, the connection is denied.
Finally, if the address is in neither file, the connection is allowed.

It is not necessary (or even possible) to list every single address that may connect to your computer. The `hosts.allow` and `hosts.deny` files enable you to specify entire subnets and groups of addresses. You can even use the keyword `ALL` to specify all possible addresses. You can also restrict specific entries in these files so they only apply to specific network services. Let's look at an example of a typical pair of `hosts.allow` and `hosts.deny` files.

```
#
# hosts.allow   This file describes the names of the hosts are
#               allowed to use the local INET services, as decided
#               by the '/usr/sbin/tcpd' server.
#

imapd, ipop3d: 199.170.177.
in.telnetd: 199.170.177., .glaci.com
ftpd: ALL
#
# hosts.deny    This file describes the names of the hosts which are
#               *not* allowed to use the local INET services, as
#               decided by the '/usr/sbin/tcpd' server.
#

ALL: ALL
```

The preceding example is a rather restrictive configuration. It allows connections to the `imap`, `ipop3d`, and `telnet` services from certain hosts, but then denies all other connections. Let's examine the files in detail.

As usual, lines beginning with a `#` character are comments and are ignored by `xinetd` when it parses the file. Each noncomment line consists of a comma-separated list of daemons followed by a colon (`:`) character and then a comma-separated list of client addresses to check. In this context, a client is any computer that attempts to access a network service on your system.

A client entry can be a numeric IP address (such as 199.170.177.25) or a host name (such as `dexter.glaci.com`) but is more often a wildcard variation that specifies an entire range of addresses. A client entry can take four different forms. The online manual page for the hosts.allow file describes

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

them as follows:

A string that begins with a dot (`.`) character. A host name is matched if the last components of its name match the specified pattern. For example, the pattern .tue.nl matches the host name wzv.win.tue.nl.

A string that ends with a dot (`.`) character. A host address is matched if its first numeric fields match the given string. For example, the pattern 131.155. matches the address of (almost) every host on the Eindhoven University network (131.155.x.x).

A string that begins with an at (`@`) sign is treated as an NIS (formerly YP) netgroup name. A host name is matched if it is a host member of the specified netgroup. Netgroup matches are not supported for daemon process names or for client user names.

An expression of the form *n.n.n.n*/*m.m.m.m* is interpreted as a *net*/*mask* pair. A host address is matched if *net* is equal to the bitwise AND of the address and the mask. For example, the net/mask pattern 131.155.72.0/255.255.254.0 matches every address in the range 131.155.72.0 through 131.155.73.255.

The example `host.allow` contains the first two types of client specification. The entry 199.170.177. will match any IP address that begins with that string, such as 199.170.177.25. The client entry `.glaci.com` will match host names such as `dexter.glaci.com` or `scooby.glaci.com`.

Let's examine what happens when a host named `daffy.glaci.com` (with IP address 199.170.179.18) connects to your Red Hat Linux box using the Telnet protocol:

1. `xinetd` receives the connection request.
2. `xinetd` begins comparing the address and name of `daffy.glaci.com` to the rules listed in `/etc/hosts.allow`. It starts at the top of the file and works its way down the file until finding a match. Both the daemon (the program handling the network service on your Red Hat Linux box) and the connecting client's IP address or name must match the information in the `hosts.allow` file. In this case, the second rule that is encountered matches the request:

   in.telnetd: 199.170.177., .glaci.com
3. Daffy is not in the 199.170.177 subnet, but it is in the `glaci.com` domain. `xinetd` stops searching the file as soon as it finds this match.

How about if Daffy connects to your box using the IMAP protocol? In this case, it matches none of the rules in `hosts.allow;` the only line that refers to the `imapd` daemon does not refer to the 199.170.179 subnet or to the `glaci.com` domain. `xinetd` continues on to the `hosts.deny` file. The entry `ALL: ALL` matches anything, so `tcpd` denies the connection.

The `ALL` wildcard was also used in the `hosts.allow` file. In this case, we are telling `xinetd` to permit absolutely any host to connect to the FTP service on the Linux box. This is appropriate for running an anonymous FTP server that anyone on the Internet can access. If you are not running an anonymous FTP site, you probably should not use the `ALL` flag.

A good rule of thumb is to make your `hosts.allow` and `hosts.deny` files as restrictive as possible and then explicitly enable only those services that you really need. Also, grant access only to those systems that really need access. Using the ALL flag to grant universal access to a particular service may be easier than typing in a long list of subnets or domains, but better a few minutes spent on proper security measures than many hours recovering from a break-in.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip You can further restrict access to services using various options within the /etc/xinetd.conf file, even to the point of limiting access to certain services to specific times of the day. Read

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Protecting Your Network with Firewalls

What is a firewall? In the non-computer world, a firewall is a physical barrier that keeps a fire from spreading. Computer firewalls serve a similar purpose, but the "fires" that they attempt to block are attacks from crackers on the Internet. In this context, a firewall, also known as a *packet filter*, is a physical piece of computer hardware that sits between your network and the Internet, regulating and controlling the flow of information.

The most common types of firewalls used today are filtering firewalls. A filtering firewall filters the traffic flowing between your network and the Internet, blocking certain things that may put your network at risk. It can limit access to and from the Internet to only specific computers on your network. It can also limit the type of communication, selectively permitting or denying various Internet services.
For Linux to act as a filtering firewall, you can use either the `ipchains` or `iptables` features. The `iptables` feature is the newer of the two and has replaced `ipchains` as the default Red Hat Linux firewall. Although ipchains has been removed from Red Hat Enterprise Linux, it is still in the Fedora Core distribution. Both ipchains and iptables are described in this chapter.

Note You can read about differences between ipchains and iptables in the Iptables HOWTO, at http://netfilter.samba.org/unreliable-guides/packet-filtering-HOWTO/packet-filtering-HOWTO.linuxdoc-7.html.

Configuring a simple firewall

You want your computer to be safe from network intruders, but you may not aspire to be a firewall expert. Instead of trying to understand all the ins and outs of `iptables` or `ipchains`, Red Hat Linux offers the Configure Firewalling window (also refers to as Lokkit).

With Configure Firewalling, you don't need to know how to load rules or memorize port numbers. All you need to do is choose what types of services you want your computer to allow in and out over your network connections. It automatically builds an iptables firewall, without requiring you to look at an iptables command.
To launch Configure Firewalling from the Red Hat menu, click System Tools → More System Tools → Lokkit. Figure 14-1 shows the Configure Firewalling window.

Figure 14-1: Build your firewall by answering a series of questions.

With Configure Firewalling open, you are ready to start. Two warnings before you begin:

The Configure Firewalling window overwrites rules created previously. Just to be safe, you should back up your previous rules as follows:
cp /etc/sysconfig/iptables /etc/sysconfig/iptables.old
Run this procedure from the console (not over the network) because if you make a mistake setting up your firewall, you might block network access and not be able to reconnect to your computer to fix the problem.

The following steps will help you answer questions about your firewall:

1. **Basic Configuration** — Choose High Security if you mostly use the computer as a

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

client for Web browsing, but not as a server (no shared printers, file systems, and so on). Low Security is best to begin with for a server computer. Only choose Disable Firewall if no connections to outside networks exist.

2. **Local Hosts** — If you are connected to a local network that does not access the Internet, you can choose to trust your connection to that network (Yes). Say No, for example, if the interface provides a connection to a cable modem or DSL line.
3. **DHCP** — Choose Yes if your network addresses are assigned via DHCP (say No if you are using static IP addresses you assign yourself).
4. **Services** — Even as a mostly-client computer, you may want to share some network services. If so, select Yes and you will be able to activate the following services:
    1. **Web Server** — If you configure Apache, TUX, or other Web server for your computer (see Chapter 21), click Yes to share that service on the network.
    2. **Incoming Mail** — If you configure Sendmail, Postfix, or other mail server on your computer (see Chapter 19), click Yes to allow your mail server to receive mail.
    3. **Secure Shell** — To allow someone to log in to your computer using a secure shell (ssh command), select Yes here.
    4. **Telnet** — To let someone log in to your computer over a telnet connection, click Yes here. (Ways of securing telnet connections are described earlier in this chapter.)
5. **Activate the Firewall** — Click Finish to activate the new firewall rules.

At this point, your new firewall is in effect. If you want to see your current firewall, type **iptables -L | less** . To see the rules used to build that firewall, open the `/etc/sysconfig/iptables` file. (For details on your iptables firewall, see the "Configuring an iptables firewall" section later in this chapter.)

Configuring an ipchains firewall

`Ipchains` works by examining packets as they are sent and received on a network interface and deciding which packets should be delivered and which should be stopped. It does this by examining a list (also called a chain) of rules. It stops at the first rule that matches the packet and examines that rule's target.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note In Red Hat Linux 8, iptables replaced ipchains as the firewall created during installation. To use ipchains, you must turn off iptables and create ipchains rules by hand. Then you must start

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To use the ipchains feature, you must have installed the ipchains RPM. To check if ipchains is installed, type **rpm -q ipchains**. To see if it's running on your system, type the **ipchains -L** command. If ipchains is not configured, you will see the message `ipchains: Incompatible with this kernel.` If ipchains is up and running, you will see output that looks like the following:

ipchains -L

Chain input (policy ACCEPT):

target prot opt source destination ports

ACCEPT all ------ anywhere anywhere n/a

ACCEPT udp ------ a.myisp.net anywhere domain - any

ACCEPT udp ------ b.myisp.net anywhere domain - any

DENY tcp -y---- anywhere anywhere any - any

DENY udp ------ anywhere anywhere any - any

Chain forward (policy ACCEPT):

Chain output (policy ACCEPT):

The preceding output represents a highly secure firewall. Notice that only input is restricted. In other words, restrictions are on which services outside users can request. The first `ACCEPT` line results from a rule that allows all requests made from users on the local system (that is, it allows the loopback driver, as indicated by n/a under ports). With that enabled, you can request any service from your local system without the packet being denied.

The next two `ACCEPT` lines allow the computers I indicate as my DNS servers (`from /etc/resolv.conf`) to request DNS services (domain) from my computer. The last two rules (`DENY`) result in all `tcp` and `udp` requests to be denied that don't match previous rules.

This firewall configuration is set in the `/etc/sysconfig/ipchains` file. When the ipchains service starts during system boot time (`/etc/init.d/ipchains`), the service reads the rules from `/etc/sysconfig/ipchains`. Here is what the rules from that file look like to create the preceding configuration:

:input ACCEPT

:forward ACCEPT

:output ACCEPT

-A input -s 0/0 -d 0/0 -i lo -j ACCEPT

-A input -s 192.160.0.253 53 -d 0/0 -p udp -j ACCEPT

-A input -s 192.160.0.254 53 -d 0/0 -p udp -j ACCEPT

-A input -s 0/0 -d 0/0 -p tcp -y -j DENY

-A input -s 0/0 -d 0/0 -p udp -j DENY

Because ipchains rules are cleared and reloaded based on the contents of this file, each time you start your system, the `/etc/sysconfig/ipchains` file is a good place to set up your firewall rules. With computer-cracker attacks on the rise, the current approach to security the experts recommend is to be secure by default. This means you should start by restricting most services and should then add only those services you want enabled. Put in the rules that allow services first, and then have all other services denied in the last two lines.

Understanding ipchains firewall rules

An ipchains target can be a simple command like `ACCEPT` or `DENY`, or it can be the name of another rule chain to begin examining. There are three default rule chains that the kernel will always examine. They are the input, output, and forward chains. You can create additional user-defined chains and call them from these original three, but for simple firewall configurations, the standard three should be sufficient.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip

ipc hai ns sto ps ex am inin g a rul e ch ain aft er find ing a ma tch , so pa y att ent ion to the ord er of rul es. Rul es wit h sp ecif ic co ndi tio ns sh oul d go bef ore tho se wit h si mil ar but bro ad er co ndi

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The general syntax is to invoke `ipchains` with a parameter specifying the action to take, followed by the rule chain to take it on. A rule description and a rule target may follow this. Table 14-3 shows action parameters you can use with `ipchains`.

Table 14-3: Ipchains Action Parameters

Action Parameter	Description
`-A, --append`	Append a new rule to the end of the specified list.
`-D, --delete`	Delete a rule from the specified list. You can specify the rule by its numeric place in the list or by the rule parameters that match it.
`-R, --replace`	Replace a rule with a new one.
`-I, --insert`	Insert a new rule into a specific position in the list.
`-L, --list`	List all the rules in a chain. If the chain name is left off, list all rules in all chains.
`-F, --flush`	Flush all the rules out of a chain.
`-Z, --zero`	Zero out the packet counters for all chains.
`-N, --new-chain`	Create a new chain with the specified name.
`-X, --delete-chain`	Delete the chain with the specified name.
`-P, --policy`	Set the policy for the chain to the specified target. The policy of a chain describes what action to take if no rule matches the packet. The default policy for all chains is ACCEPT.
`-M, --masquerading`	Allows viewing of masqueraded connections. IP Masquerading is discussed in Chapter 16.
`-S, --set`	Set the timeouts for TCP, TCPFIN, and UDP packets.
`-C, --check`	Check a supplied packet against the given chain. This is useful mainly for debugging.
`-h`	Print a Help message describing parameters to ipchains.

As you can see in Table 14-3, the ipchains action parameters can be expressed in two forms, either as a dash followed by a single capital letter, or two dashes followed by a descriptive word. Both will work, so use whichever you prefer. In the example in this chapter, I will use the abbreviated version.

Usually, we follow the action parameter with the rule chain to apply it to. Rules added to the input chain will be examined only when filtering network packets being received by the Linux box. Similarly, the output chain is examined only for packets being transmitted from the Linux box. The forward chain is examined only for network packets that are received by the Linux system but will be delivered to some other network system. Packet forwarding only occurs when your system is configured as a router. After specifying a chain to act on, you may specify some optional parameters to define a rule. Table 14-4 lists the available optional parameters.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 14-4: Ipchains Optional Parameters

Parameter	Description
`-p, --protocol[!] protocol`	Specify the protocol that the rule should match against. This should be TCP, UDP, or ICMP.
`-s, --source [!] address`	Specify the source address to match against. This can be an individual address, or you can specify an entire subnet by following the address with a / and the number of 1 bits in the left side of the subnet mask. Thus, the address 199.170.177.0/24 would have a subnet mask of 255.255.255.0.
`--source-port [!]port`	The source TCP or UDP port number as specified in `/etc/services`. You can also specify a range of ports by listing the first and last port number separated by a ':' colon character.
`-d, --destination [!] address`	Specify the destination address to match against. This can be an individual address, or you can specify an entire subnet by following the address with a / and the number of 1 bits in the left side of the subnet mask. Thus, the address 199.170.177.0/24 would have a subnet mask of 255.255.255.0.
`--destination-port [!] port`	The destination TCP or UDP port number as specified in `/etc/services`. You can also specify a range of ports by listing the first and last port number separated by a ':' colon character.
`--icmp-type [!] typename`	Set the type of ICMP packet to use.
`-j, --jump target`	The name of the target (action) to execute when the rule matches. This could be the name of another ipchain or one of several predefined targets.
`-i, --interface [!] name`	The name of the network interface that this rule applies to. If this option is not supplied, the rule will apply to all interfaces.
`[!] -f, --fragment`	The rule will apply only to fragmented packets, excluding the first packet. In other words, it applies to all packet fragments after the first one.
`-b, --bidirectional`	The rule should apply to both incoming and outgoing packets.
`-v, --verbose`	Print debug messages when processing this ipchains command.
`-n, --numeric`	Use IP addresses instead of host names when printing output to the screen.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 14-4: Ipchains Optional Parameters

Parameter	Description
`-l, --log`	Turn on kernel logging of matching packets. This will slow things down and fill up your hard drive. It is intended mainly for debugging.
`-o, --output [maxsize]`	Divert packets to a user space process. Another debugging feature.
`-m, --mark markvalue`	Mark the packet with a 32-bit signature. This is probably only useful to you if you are a kernel hacker.
`-t, --TOS andmask xormask`	Examine the TOS field of the packet using the supplied bit masks. Read the ipchains man page for a complete discussion of this option.
`-x, --exact`	Display exact values of packet counters rather than numbers rounded to the kilobyte.
`[!] -y, --syn`	Examine the SYN bit in the TCP packet being looked at. Useful for blocking TCP connection from being initiated in one direction but not the other.
`--line-numbers`	Show line numbers when listing rules. This is useful if you plan to delete or modify rules by position number.
`--no-warnings`	Disable all warning messages.

After specifying a rule for a particular type of packet, you must specify the target for it using the `-j` or `--jump` option. This tells ipchains what to do with that packet when it finds a rule that matches it. The target could be the name of another rule chain to traverse, but more often it is one of the predefined actions described in Table 14-5.

Table 14-5: ipchains Targets

Target	Description
`ACCEPT`	Accept the packet and deliver it in the normal way.
`DENY`	Drop the packet completely.
`REJECT`	Drop the packet and then send an ICMP packet with an explanation to the sending host. This is primarily useful for debugging.
`MASQ`	Use IP Masquerading for this packet type. Refer to Chapter 16 for an explanation of masquerading.
`REDIRECT`	Redirect the packet to a new location.
`RETURN`	Return from this chain to the chain that called it. Continue examining rules in the calling chain where you left off.

I've shown you the various components of an `ipchains` command. It is time to put them together into

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

some practical examples. It is possible to create some very sophisticated and complicated rule lists with `ipchains`, but I will keep my examples rather simple. Keeping things simple is generally a good policy, since large, complicated rule chains can impact system performance. More time spent examining rules means less time delivering packets and serving up information. The higher the traffic level on your Linux box, the greater the performance impact of those complicated rule chains.

Changing ipchains firewall rules

Now let's try adding a rule. As an example, imagine we want to block ICMP packets to disallow "pinging" of our Linux box. You may do that to avoid various denial-of-service attacks that could be launched against your system. Block ICMP with a command like the following:

ipchains -A input -p icmp -j DENY
This specifies that we are adding a rule to the input chain. It will match any ICMP packet and will drop it rather than allowing it through. Now if you are using the `ping` command against your Linux box, you should receive no response. Type the **ipchains -L** command again, and you'll see something like this:

Chain input (policy ACCEPT):

target prot opt source destination ports

DENY icmp ------ anywhere anywhere any -> any

Chain forward (policy ACCEPT):

Chain output (policy ACCEPT):
You can see your new rule listed. This rule will block all ICMP packets entering your system, regardless of which computer sent those packets. If your Linux system is acting as a router, it will also block ICMP packets that are being forwarded from the Internet to your network, or vice versa. People on the Internet will be unable to ping anything on your network. Likewise, you will be unable to ping anything on the Internet. Perhaps that is not what you want. Assume then that you wish to block pinging of systems on your network by people on the Internet, but allow pinging of the router and allow the router to ping hosts on the Internet. First, we should flush the contents of the input chains using the `-F` parameter; then we can add our new rule.

ipchains -F input

ipchains -A forward -p icmp -j DENY
Now we can ping the Linux system and the Linux system can ping other boxes, but ping requests will not be passed through the Linux system. If you wish, use the `ipchains -L` command to verify that the rule has now been added to the forward chain rather than the input chain.

You may also wish to block the telnet protocol when coming from the Internet. For this example, let us assume that our Linux router is connected to the Internet via a dial-up connection called ppp0 and is connected to our internal LAN via an Ethernet connection called eth0. In that case, you could block telnet with a command like the following:

ipchains -A input -i ppp0 -p tcp --dport 23 -j DENY
This rule basically says that any TCP packet with a destination port of 23 (the telnet port as specified in `/etc/services`) that is arriving on the ppp0 interface should be dropped. This does not prevent you from telneting to your Linux box from your internal network, but it does block telnet access from the Internet.
I'm going to finish up with one more useful example. Imagine you want to allow any type of outbound TCP connection to the Internet, but want to block any inbound TCP connection. Every TCP connection sends packets in both directions, so at first glance it would seem impossible. Block all inbound TCP packets, and the reply packets to your outbound connections will also be blocked. The trick is to block only the initial TCP packet that is used to start an inbound connection. We can do this because all TCP connection requests start with a packet that has something called the SYN bit set. We can use the `--syn` option to tell ipchains to look for that bit. Try the following command:

ipchains -A input -i ppp0 -p tcp --syn -j DENY
There are many useful ways to filter traffic using ipchains. I encourage you to read the ipchains man page (type **`man ipchains`**) and the ipchains HOWTO document to learn more about it. You can find the HOWTO document at www.tldp.org/HOWTO/IPCHAINS-HOWTO.html.

Saving ipchains firewall rules

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

After you have created the ipchains rules you want, it is important to save them to a file; otherwise, they will be lost when you reboot the server. Fortunately, a pair of useful scripts (`ipchains-save` and `ipchains-restore`) is provided for exactly this purpose. Essentially, `ipchains-save` will echo the current ipchains rule list to the screen. The `ipchains-restore` script reads in the specially formatted rule list and makes it active. After customizing `ipchains`, save the new rules to a file by running `ipchains-save` and directing the output to a file such as:

/sbin/ipchains-save > /tmp/ipchains.rules

Next, you can add the rules you have created to the `/etc/sysconfig/ipchains` file. As root user, open the file in any text editor. Then read in the `ipchains.rules` file you have created in the preceding example.

Configuring an iptables firewall

The iptables firewall feature (also referred to as netfilter) is the default firewall software when you install Red Hat Linux. Today, both `ipchains` and `iptables` are available with Red Hat Linux, although only one can be active at a time. `iptables` is the better supported firewall tool these days.

If you currently use `ipchains` on your Red Hat Linux system and you wish to convert to `iptables`, you need to disable `ipchains` before you can use `iptables`. In fact, if you try to use the `iptables` script (`/etc/init.d/iptables`) to start `iptables` while the `ipchains` module is loaded, the `iptables` startup script silently exits.

While `iptables` is generally considered to be more complex than `ipchains` to work with, it is also considered to be more powerful and flexible. Because development is not continuing for `ipchains`, the few broken `ipchains` features that exist will probably stay broken.

This section describes how to turn on iptables and set up firewall rules for several different types of situations. It also tells how to turn on features related to firewalls that allow your `iptables` firewall to do Network Address Translation (NAT), IP masquerading, port forwarding, and transparent proxies.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The Red Hat Linux startup scripts will "punch a hole" through your firewall if you use certain services, and will therefore work even if they are not explicitly enabled in your iptables

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Turning on iptables

To turn on `iptables`, you might have to turn off `ipchains`. (This might be the case if you just upgraded from a release prior to Red Hat Linux 8, where `ipchains` was the default.) Then you can add the modules and create the rules you need for your `iptables` firewall. The following procedure describes how to get `iptables` going on your Red Hat Linux system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip During the time that you turn off ipc hains, turn on ipt ables, and config ure the filte ring rul es for ipt abl es, your comp ute r may be un pro tec ted fro m the net wor k. You should test your fire

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

1. Stop the ipchains script from starting automatically at boot time:
2. # **chkconfig ipchains off**
3. Set the iptables script to start automatically at boot time:
4. # **chkconfig iptables on**
5. Now you can either reboot Red Hat Linux, or stop ipchains and unload the ipchains module, as follows:
6. # **service ipchains stop**
7. # **modprobe -r ipchains**
8. Before you can start iptables you must have a working set of rules that have been placed in your `/etc/sysconfig/iptables` file. To create those rules, refer to the examples in the following sections. (Without the configuration file in place, iptables fails silently.)
9. If you are doing NAT or IP Masquerading, turn on IP packet forwarding. One way to do this is to change the value of `net.ipv4.ip_forward` to 1 in the `/etc/sysctl.conf` file. Open that file as root user with any text editor and change the line to appear as follows:

   net.ipv4.ip_forward = 1
10. Restart your network interfaces to have IP packet forwarding take effect.
11. # **/etc/init.d/network restart**
12. Once the rules file is in place, start up iptables:
13. # **/etc/init.d/iptables start**
14. At this point, iptables is installed as your firewall. You can check to see that the modules used by iptables are loaded by using the `lsmod` command, as follows:
15. # **lsmod |grep ip**
16. ipt_REJECT 3928 6 (autoclean)
17. iptable_filter 2412 1 (autoclean)
    ip_tables 15096 2 [ipt_REJECT iptable_filter]
18. If you want to allow passive FTP or IRC connections from computers on your LAN, you may need to load those modules by adding them to the `/etc/modules.conf` file. (See the description of passive FTP and IRC after the firewall-example sections.)

If your iptables service didn't start, make sure that:

An ipchains module is not loaded. (Unload it with `modprobe -r` as shown previously.)

The `/etc/sysconfig/iptables` file exists. (You need to create one if one was not already created for you when you installed Red Hat Linux.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip As you add ipt ables rules, more modules will have to be loaded. Appropriate modules should be loaded automatically when a new rule is entered. Run lsmod | grep ip again after ye

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The following sections contain examples of iptables firewall rules.

Creating iptables firewall rules

One way to configure iptables is to start by adding and deleting rules to your kernel from the command line. Then when you get a set of rules that you like, you save the rules that are currently running on your system. The tools you use to create your firewall rules and then make them permanent are as follows:

1. **`iptables`** — Use this command to append (-A), delete (-D), replace (-R) or insert (-I) a rule. Use the -L option to list all current rules.
2. **`iptables-save -c`** — Use this command to save the rules from the kernel and install them in the configuration file.
3. **`/etc/sysconfig/iptables`** — This is the configuration file that contains the rules that were saved from the `iptables-save` command.
4. **`/etc/init.d/iptables`** — This is the iptables start-up script that must run automatically each time Red Hat Linux reboots. When it starts, it clears all iptables rules and counters and installs the new rules from the `/etc/sysconfig/iptables` file. You can also use this script with different options from the command line to check the status of iptables (`status`) or to run `iptables-save` for you to save the current rules (`save`).

To get you started with iptables, I'm providing a sample set of iptables rules along with descriptions of how you might change those rules for different situations. Here's how you could load and save any of the sets of rules described in the following example:

If you are currently running ipchains, complete the procedure in the previous "Turning on iptables" section.
Stop iptables and clear all existing rules:
/etc/init.d/iptables stop
Add the rules shown in the following example to a file, using any text editor. Modify the rules to suit your situation and save the file.
As root user, run the file as a shell script. For example, if you named the file `firescript`, you could run it as follows:
sh firescript
See how the rules were loaded into the kernel:
iptables -L
If everything looks okay, save the rules that are now in the kernel into the `/etc/sysconfig/iptables` file:
iptables-save > /etc/sysconfig/iptables

From now on the rules will be read each time you reboot or restart iptables. Save a copy of the script you used to create the rules, in case you ever need it again.

Example 1: Firewall for shared Internet connection (plus servers)

This example features a home or small-office LAN with a Red Hat Linux system acting as an iptables firewall between the LAN and the Internet. The firewall computer also acts as a Web server, FTP server, and DNS server. Figure 14-2 shows this configuration.

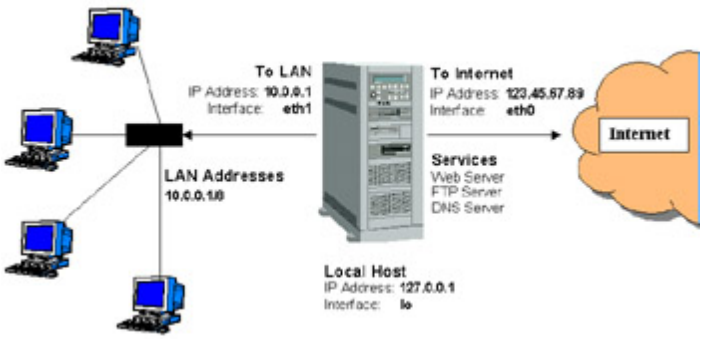

Figure 14-2: Using iptables as a firewall between the Internet and a LAN.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

If you want to use the sample firewall script that follows, you must change the following information to match your configuration:

Firewall computer — The firewall computer is set up as follows:

Local host — `127.0.0.1` (IP address) and `lo` (interface). You shouldn't need to change these.

Connection to the Internet — `123.45.67.89` (IP address) and `eth0` (interface). Replace them with the static IP address and interface name associated with your connection to the Internet, respectively.

Connection to the LAN — `10.0.0.1` (IP address) and `eth1` (interface). Replace `10.0.0.1` and `eth1` with the static IP address and interface name associated with your connection to your LAN, respectively.

Computers on the LAN — Each computer on the LAN in the example has an IP address between `10.0.0.2` to `10.0.0.254`. Change `10.0.0.255` to a number that matches your LAN's range of addresses.

Here is an example of a script to load firewall rules that could be used for the configuration shown in Figure 14-2:

```
# (1) Policies (default)
iptables -P INPUT DROP
iptables -P OUTPUT DROP
iptables -P FORWARD DROP

# (2) User-defined chain for ACCEPTed TCP packets
iptables -N okay
iptables -A okay -p TCP --syn -j ACCEPT
iptables -A okay -p TCP -m state --state ESTABLISHED,RELATED -j ACCEPT
iptables -A okay -p TCP -j DROP

# (3) INPUT chain rules

# Rules for incoming packets from LAN
iptables -A INPUT -p ALL -i eth1 -s 10.0.0.0/8 -j ACCEPT
iptables -A INPUT -p ALL -i lo -s 127.0.0.1 -j ACCEPT
iptables -A INPUT -p ALL -i lo -s 10.0.0.1 -j ACCEPT
iptables -A INPUT -p ALL -i lo -s 123.45.67.89 -j ACCEPT
iptables -A INPUT -p ALL -i eth1 -d 10.0.0.255 -j ACCEPT

# Rules for incoming packets from the Internet

# Packets for established connections
iptables -A INPUT -p ALL -d 123.45.67.89 -m state --state
ESTABLISHED,RELATED -j ACCEPT

# TCP rules
iptables -A INPUT -p TCP -i eth0 -s 0/0 --destination-port 21 -j okay
iptables -A INPUT -p TCP -i eth0 -s 0/0 --destination-port 22 -j okay
iptables -A INPUT -p TCP -i eth0 -s 0/0 --destination-port 80 -j okay
iptables -A INPUT -p TCP -i eth0 -s 0/0 --destination-port 113 -j okay
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
# UDP rules
iptables -A INPUT -p UDP -i eth0 -s 0/0 --destination-port 53 -j ACCEPT
iptables -A INPUT -p UDP -i eth0 -s 0/0 --destination-port 2074 -j ACCEPT
iptables -A INPUT -p UDP -i eth0 -s 0/0 --destination-port 4000 -j ACCEPT

# ICMP rules
iptables -A INPUT -p ICMP -i eth0 -s 0/0 --icmp-type 8 -j ACCEPT
iptables -A INPUT -p ICMP -i eth0 -s 0/0 --icmp-type 11 -j ACCEPT

# (4) FORWARD chain rules
# Accept the packets we want to forward
iptables -A FORWARD -i eth1 -j ACCEPT
iptables -A FORWARD -m state --state ESTABLISHED,RELATED -j ACCEPT

# (5) OUTPUT chain rules
# Only output packets with local addresses (no spoofing)
iptables -A OUTPUT -p ALL -s 127.0.0.1 -j ACCEPT
iptables -A OUTPUT -p ALL -s 10.0.0.1 -j ACCEPT
iptables -A OUTPUT -p ALL -s 123.45.67.89 -j ACCEPT

# (6) POSTROUTING chain rules
iptables -t nat -A POSTROUTING -o eth0 -j SNAT --to-source 123.45.67.89
```

I divided the commands in the preceding script into six sections. The following text describes each of those sections.

(1) Policies — The `iptables -P` commands set the default policies for `INPUT`, `OUTPUT`, and `FORWARD` chains. By assigning each of those policies to `DROP`, any packet that isn't matched is discarded. In other words, for a packet to get through, it has to be specifically matched and ACCEPTed by one of the other rules in the script.

(2) User-defined chain — A user-defined chain I call `okay` is created to do a few more checks on packets requesting certain TCP services that I'm going to allow through (Web, FTP, and DNS services). The `-N okay` option creates the `okay` chain. The next line says that a SYN packet (`--syn`), which requests a new connection, is fine to let through. The next line allows through packets associated with an `ESTABLISHED` connection (one that has already had traffic pass through the interface) or a `RELATED` connection (one that is starting a new connection related to an already-established connection). The final line in this set tells iptables to `DROP` packets that don't meet one of those checks.

(3) INPUT chain rules — The bulk of the packet filtering is done in the `INPUT` chain. The first sets of input rules indicates to iptables when to always accept packets from the Internet and from the LAN. The next three sets determine which requests for specific protocols (`TCP`, `UDP`, and `ICMP`) are accepted.

> **Packets from LAN** — Because you want the users on your LAN and the firewall computer itself to be able to use the Internet, this set of rules lets through packets that are initiated from those computers. The first line tells iptables to accept packets for `ALL` protocols for which the source is an acceptable addresses on your LAN (`-s 10.0.0.0/8`, which represents IP numbers 10.0.0.1 through 10.0.0.254). The next three lines allow packets that come from all valid IP addresses on the firewall computer itself (`-s 127.0.0.1`, `10.0.0.1`, and `123.45.67.89`). The last line accepts broadcast packets (`-d 10.0.0.255`) on the LAN.
>
> **Packets from Internet (already connected)** — This line is split in two (I just used the backslash to join the lines because the page wasn't wide enough to

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

show it as one line). It `ACCEPT`s packets that are both associated with connections that are already established (`--state ESTABLISHED,RELATED`) and are requested directly to the firewall's IP address (`123.45.67.89`).

TCP rules (new connections) — Here is where you open up the ports for the TCP services you want to provide to anyone from the Internet. In these lines you open ports for FTP service (`--destination-port 21`), secure shell service (`22`), Web service (`80`), and `IDENTD` authentication (`113`), the last of which might be necessary for protocols such as IRC. Instead of accepting these requests immediately, you jump to the `okay` chain you defined to further check that the packets was formed properly.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution You want to ensure that the services on the ports to which you are allowing access are properly configured before you allow packets to be accepted to them.

UDP rules (new connections) — These lines define the ports where connection packets are accepted from the Internet for UDP services. In this example I chose to accept requests for DNS service (`--destination-port 53`) because the computer is set up as a DNS server. The example also

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

illustrates lines that accept requests for a couple of other optional ports. Port 2074 is needed by some multimedia applications the users on your LAN might want to use, and port 4000 is used by the ICQ protocol (for online chats).
ICMP rules — ICMP messages are really more for reporting conditions of the server than for actually providing services. Packets from the Internet that are accepted for ICMP protocol requests are those for ICMP types 8 and 11. Type 8 service, which allows your computer to accept *echo reply messages*, makes it possible for people to `ping` your computer to see if it is available. Type 11 service relates to packets whose *time to live* (TTL) was exceeded in transit, and for which you are accepting a Time Exceeded message that is being returned to us. (You need to accept Type 11 messages to use the `traceroute` command to find broken routes to hosts you want to reach.)

(4) FORWARD chain rules — Because this firewall is also acting as a router, `FORWARD` rules are needed to limit what the firewall will and will not pass between the two networks (Internet and local LAN). The first line forwards everything from the local LAN (`-A FORWARD -i eth1`). The second line forwards anything from the Internet that is associated with an established connection (`--state ESTABLISHED, RELATED`).
(5) OUTPUT chain rules — These rules basically exist to prevent anyone from your local computer from *spoofing* IP addresses (that is, from saying packets are coming from somewhere that they are not). According to these three rules, each packet output from your firewall must have as its source address one of the addresses from the firewall computer's interfaces (127.0.0.1, 10.0.0.0.1, or 123.45.67.89).
(6) POSTROUTING chain rules — The POSTROUTING chain defines rules for packets that have been accepted, but need additional processing. This is where the actual network-address translation (NAT) work takes place. For the NAT table (`-t nat`), in the `POSTROUTING` chain, all packets that go out to the Internet have their addresses translated to that of the firewall's external interface (`--to-source 123.45.67.89`). In this case I used the Source Network Address Translation (`SNAT`) chain because I have a static IP address (`123.45.67.89`) associated with my Internet connection. If I were using a dynamic IP address (via DHCP), I would use `MASQUERADE` instead of `SNAT`. I would also have to change any references to `-d 123.45.67.89` to `-i eth0`. (Of course, you would be using a different IP address and possibly a different Ethernet interface.)

Example 2: Firewall for shared Internet connection (no servers)

In this scenario, the firewall is protecting a Linux system (firewall) and several systems on the LAN that only want to connect to the Internet. No servers are behind this firewall, so you want to prevent people from the Internet from requesting services.
You could use the same script shown in Example 1, but not use lines that request TCP services. So you could drop the user-defined chain in section 2 and drop the TCP rules from section 3. You could also remove the ICMP rules, if you don't want your firewall to be invisible to `ping` requests (Type 8) and if you don't care about receiving messages when your packets exceed their time to live values (Type 11), such as when a packet runs into a broken router.

Example 3: Firewall for single Linux system with Internet connection

In this example, there is one network interface, connecting your Red Hat Linux system to the Internet. You are not sharing that connection with other computers and you are not offering any services from your computer.
In this case, you could cut sections 2, 4, and 6. From section 3, you could cut all rules relating to incoming requests from the LAN and all TCP services. As I mentioned in Example 2, you could also remove the Type 8 ICMP rule to make your firewall invisible to `ping` requests from the Internet and the Type 11 ICMP rule to not accept messages about failed time-to-live packets.

Understanding iptables

Now that you've seen something about what iptables rules look like and how you can get them going, step back a bit and see how iptables works.

The iptables feature works by having IP packets (that is, network data) that enter or leave the firewall computer, traverse a set of chains that define what is done with the packet. Each rule that you add

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

essentially does the following:

Checks if a particular criterion is met (such as that a packet requests a particular service or comes from a particular address) and

Takes an action (such as dropping, accepting, or further processing a packet).

Different sets of rules are implemented for different types of tables. For example, tables exist for filtering (`filter`), network address translation (`nat`), and changing packet headers (`mangle`). Depending on the packet's source and destination, it traverses different chains. Most of the rules you create will relate to the `filter` table (which is implied if no other table is given).

The chains associated with the filter table are `INPUT`, `OUTPUT`, and `FORWARD`. You can add user-defined chains as well. You will probably be most interested in adding or removing particular TCP and UDP services using the basic rules shown in the previous example. Assign `ACCEPT` to packets you want to go through and `DROP` to those you want to discard. You can also assign `REJECT` to drop a packet but return a message to the sender, or `LOG` to neither drop nor accept the message, but to log information about the packet.

A lot of great features are built into iptables. The following descriptions tell you about some cool things you can do with iptables and give you some tips on using it.

Allowing FTP and IRC services through an iptables firewall

With passive FTP, the FTP client sends its IP address and the port number on which it will listen for data to the server. If that client is on a computer that is behind your firewall, for which you are doing NAT, that information must be translated as well or the FTP server will not be able to communicate with the client.

iptables uses connection-tracking modules to track connections. Using this feature, it can look inside the FTP data themselves (that is, not in the IP packet header), to get the information it needs to do NAT (remember that computers from the Internet can't talk directly to your private IP addresses). To do FTP connection tracking (to allow passive FTP connections to the clients on your LAN), you need to have the following modules loaded:

1. `ip_conntrack`
2. `ip_conntrack_ftp`
3. `ip_nat_ftp`

The same is true for chats (IRC) and DCC sends. Addresses and port numbers are stored within the IRC protocol packets, so those packets must be translated too. To allow clients on your LAN to use IRC services, you need to load the following modules:

```
ip_conntrack_irc
ip_nat_irc
```

The default port for IRC connections is 6667. If you don't want to use the default you can add different port numbers when you load the connection-tracking modules:

insmod ip_conntrack_irc.o ports=6668,6669

Using iptables to do NAT or IP Masquerading

As noted in the iptables example, you can use Source Network Address Translation (`SNAT`) or IP Masquerading (`MASQUERADE`) to allow computers on your LAN with private IP addresses to access the Internet through your iptables firewall. Choose `SNAT` if you have a static IP address for your Internet connection and `MASQUERADE` if the IP address is assigned dynamically.

When you create the `MASQUERADE` or `SNAT` rule, it is added to the `NAT` table and the `POSTROUTING` chain. For `MASQUERADE` you must provide the name of the interface (such as `eth0`, `ppp0`, or `slip0`) to identify the route to the Internet or other outside network. For `SNAT` you must also identify the actual IP address of the interface. Here is an example of a `MASQUERADE` rule:

iptables –t nat –A POSTROUTING –o eth0 –j MASQUERADE

Here is an example of a SNAT rule:

iptables -t nat -A POSTROUTING -o eth0 -j SNAT --to-source 12.12.12.12

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You can add several source addresses if you have multiple addresses that provide a route to the Internet (for example, `--to-source 12.12.12.12.1-12.12.12.12.254`). Although `MASQUERADE` uses some additional overhead, you probably need to use it instead of `SNAT` if you have a dial-up connection to the Internet for which the IP address changes on each connection.

Remember that you need to make sure that IP forwarding is turned on in the kernel. (It is off by default.) To turn it on temporarily, you can do the following:

echo 1 > /proc/sys/net/ipv4/ip_forward

If you require dynamic IP addressing, turn on that service:

echo 1 > /proc/sys/net/ipv4/ip_dynaddr

Using iptables as a transparent proxy

With the `REDIRECT` target you can cause traffic for a specific port on the firewall computer to be directed to a different one. Using this feature you can direct host computers on your local LAN to a proxy service on your firewall computer without those hosts knowing it.
The following is an example of a set of command-line options to the `iptables` command that causes a request for Web service (port 80) to be directed to a proxy service (port 3128):

-t nat -A PREROUTING -p tcp --dport 80 -j REDIRECT --to-ports 3128
You can only use `REDIRECT` targets in `PREROUTING` and `OUTPUT` chains within a `nat` table. You can also give a range of port numbers to spread the redirection across multiple port numbers.

Using iptables to do port forwarding

What if you have only one static IP address, but you want to use a computer other than your firewall computer to provide Web, FTP, DNS, or some other service? You can use the Dynamic Network Address Translation (`DNAT`) feature to direct traffic for a particular port on your firewall to another computer.
For example, if you wanted all requests for Web service (port 80) that were directed to the firewall computer (`-d 15.15.15.15`) to be directed to another computer on your LAN (such as 10.0.0.25), you could use the following `iptables` command:

**iptables -t nat -A PREROUTING -p tcp -d 15.15.15.15 -dport 80 **

-j DNAT --to-destination 10.0.0.25

(Note that the preceding example should actually appear on one line. The backslash indicates continuation on the next line.)
You can also spread the load for the service you are forwarding by providing a range of IP addresses (for example, `--to-destination 10.0.0.1-10.0.0.25`). Likewise, you can direct the request to a range of ports as well.

Using logging with iptables

Using the `LOG` target you can log information about packets that meet the criteria you choose. In particular you might want to use this feature to log packets that seem like they might be improper in some way. In other words, if you don't want to drop a packet for some reason, you can just log its activity and decide later if something needs to be corrected.
The `LOG` target directs log information to the standard tools used to do logging in Red Hat Linux: `dmesg` and `syslogd`. Here is an example of a rule using a `LOG` target:

-A FORWARD -p tcp -j LOG --log-level info
Instead of `info`, you could use any of the following log levels available with `syslog`: `emerg`, `alert`, `crit`, `err`, `warning`, `notice`, `info`, or `debug`. Using the `--log` prefix option as follows, you could also add information to the front of all messages produced from this logging action:

-A FORWARD -p tcp -j LOG --log-level info --log-prefix "Forward INFO "

Enhancing your iptables firewall

You can modify or expand on the iptables examples given in this chapter in many ways. iptables is tremendously flexible.
When you actually create your own iptables firewall, you should refer to the iptables man page (type **man iptables**) for detailed descriptions of options, ways of matching, ways of entering addresses, and

other details. I also recommend an excellent iptables Tutorial by Oskar Andreasson (`http://iptables-tutorial.frozentux.net`).

Here are a few tips for using iptables features:

Reduce rules — Try to improve performance by reducing the number of rules. Using subchains can keep a packet from seeing rules that don't apply to it.

Deal with fragments — Use the `-f` option to refer to the second and subsequent packets of a packet that was split into fragments. In general, it is safe to not drop second and third fragments for which you don't have a first packet fragment because they won't be reassembled. If you use NAT the fragments are assembled before filtering, so you shouldn't have problems with unfiltered fragments being sent through.

Opposite — To make a rule its opposite, use an exclamation mark (`!`).

All interfaces — To match all interfaces of a type, use a plus sign (+), as in `eth+`.

Blocking connections — Use the `--syn` option to block SYN packets (that is, those packets requesting connections). This option only applies to TCP packets.

Limiting — Use the `--limit` option to restrict the rate of matches that result in log messages. This option allows matches to only produce messages a limited number of times per second (the default is three per hour with burst of five).

DOS Attacks — You can use the `--limit` option to reduce the impact of DOS attacks, but it still won't stop them altogether. As long as the traffic is directed at your server, your network bandwidth is being leeched away, and the machine still utilizes resources to ignore the data.

Table types — The default table type is `filter`. The other types of tables are `nat` (for IP Masquerading) and `mangle` (for altering packets). To use a table other than `filter` you must add a `-t` *`table_type`* option, where *`table_tape`* is either `nat` or `mangle`.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Detecting Intrusions from Log Files

Preparing your system for a cracker attack is only part of the battle. You must also recognize a cracker attack when it is occurring. Understanding the various log files in which Red Hat Linux records important events is critical to this goal. The log files for your Red Hat Linux system can be found in the `/var/log` directory.

Red Hat Linux comes with a System Logs window (`redhat-logview` command) that you can use to view and search critical system log files from the GUI. To open the System Logs window, from the main desktop menu, click System Tools → System Logs. Figure 14-3 shows an example of the System Logs window.

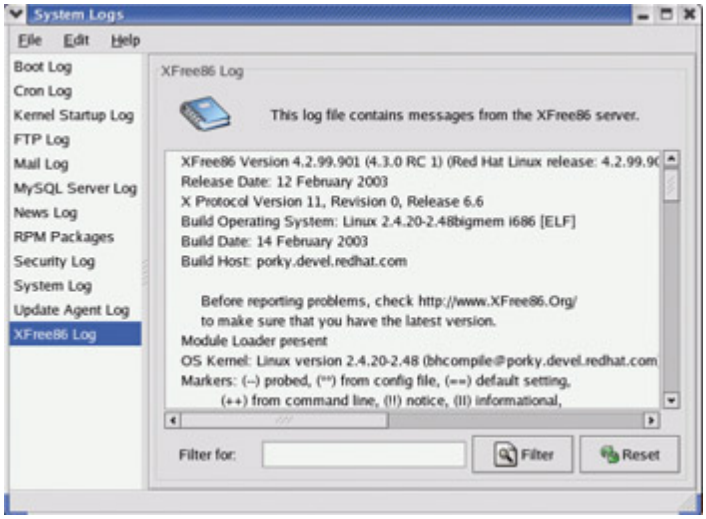

Figure 14-3: Display system log files in the System Logs window.

To view a particular log file, click the log name in the left column. If you are looking for a particular message or problem, type a keyword into the "Filter for" box, and click Filter. Only lines containing that word are displayed. Case matters, so searching for "Mem" won't find "mem" when you use the filter. Click Reset to display the whole file again.

Table 14-6 contains a listing of log files displayed on the System Logs window, along with other files in the `/var/log` directory that may interest you. (* Indicates a log file that is not contained in the System Logs window. Access these files directly from `/var/log`.)

Table 14-6: Log Files in the /var/log Directory

System Logs Name	Filename	Description
`Boot Log`	`boot.log`	Contains messages indicating which systems services have started and shut down successfully and which (if any) have failed to start or stop.
`Cron Log`	`cron`	Contains status messages from the `crond`, a daemon that periodically runs scheduled jobs, such as backups and log file rotation.
`Kernel Startup Log`	`dmesg`	A recording of messages printed by the kernel when the system boots.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 14-6: Log Files in the /var/log Directory

System Logs Name	Filename	Description
FTP Log	xferlog	Information about files transferred using the wu-ftpd FTP service.
Apache Access Log	httpd/access_log	Logs requests for information from your Apache Web server.
Apache Error Log	httpd/error_log	Logs errors encountered from clients trying to access data on your Apache Web server.
Mail Log	maillog	Contains information about addresses to which and from which e-mail was sent. Useful for detecting spamming.
MySQL Server Log	mysqld.log	Includes information related to activities of the MySQL database server (mysqld).
News Log	spooler	Directory containing logs of messages from the Usenet News server, if you are running one.
RPM Packages	rpmpkgs	Contains a listing of RPM packages that are installed on your system.
Security Log	secure	Records the date, time, and duration of login attempts and sessions.
System Log	messages	A general-purpose log file to which many programs record messages.
Update Agent Log	up2date	Contains messages resulting from actions by the Red Hat Update Agent.
XFree86 Log	XFree86.0.log	Includes messages output by the Xfree86 server.
*	gdm/:0.log	Holds messages related to the login screen (GNOME display manager).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 14-6: Log Files in the /var/log Directory

System Logs Name	Filename	Description
*	samba/log.smbd	Messages from the Samba SMB file service daemon.
*	squid/access.log	Contains messages related to the squid proxy/caching server.
*	vsftpd.log	Contains messages relating to transfers made using the vsFTPd daemon (FTP server).
*	sendmail	Error messages recorded by the sendmail daemon.
*	uucp	Status messages from the Unix to Unix Copy Protocol daemon.

The role of syslogd

Most of the files in the /var/log directory are maintained by the syslogd service. The syslogd daemon is the System Logging Daemon. It accepts log messages from a variety of other programs and writes them to the appropriate log files. This is better than having every program write directly to its own log file because it allows you to centrally manage how log files are handled. It is possible to configure syslogd to record varying levels of detail in the log files. It can be told to ignore all but the most critical message, or it can record every detail.

The syslogd daemon can even accept messages from other computers on your network. This is particularly handy because it enables you to centralize the management and reviewing of the log files from many systems on your network. There is also a major security benefit to this practice. If a system on your network is broken into, the cracker cannot delete or modify the log files because those files are stored on a separate computer. It is important to remember, though, that those log messages are not, by default, encrypted. Anyone tapping into your local network will be able to eavesdrop on those messages as they pass from one machine to another. Also, though the cracker may not be able to change old log messages, he will be able to affect the system such that any new log messages should not be trusted.

It is not uncommon to run a dedicated loghost, a computer that serves no other purpose than to record log messages from other computers on the network. Because this system runs no other services, it is unlikely that it will be broken into. This makes it nearly impossible for a cracker to erase his or her tracks. It does not, however, mean that all of the log messages are accurate after a cracker has broken into a machine on your network.

Redirecting logs to a loghost with syslogd

To redirect your computer's log files to another computer's syslogd, you must make some changes to your local syslogd's configuration file. The file that you need to work with is /etc/syslog.conf. Become root using the su command and then load the /etc/syslog.conf file in a text editor (such as vi). You should see something similar to this:

```
# Log all kernel messages to the console.
# Logging much else clutters up the screen.
#kern.*                         /dev/console


# Log anything (except mail) of level info or higher.
# Don't log private authentication messages!
*.info;mail.none;news.none;authpriv.none;cron.none  /var/log/messages
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

```
# The authpriv file has restricted access.
authpriv.*                      /var/log/secure

# Log all the mail messages in one place.
mail.*                          /var/log/maillog

# Log cron stuff
cron.*                          /var/log/cron

# Everybody gets emergency messages
*.emerg                         *

# Save news errors of level crit and higher in a special file.
uucp,news.crit                  /var/log/spooler

# Save boot messages also to boot.log
local7.*                        /var/log/boot.log

#
# INN
#
news.=crit                      /var/log/news/news.crit
news.=err                       /var/log/news/news.err
news.notice                     /var/log/news/news.notice
```

The lines beginning with a `#` character are comments. Other lines contain two columns of information, separated by colons (spaces won't work). The left field is a semicolon-separated list of message types and message priorities. The right field is the log file to which those messages should be written. To send the messages to another computer (the loghost) instead of a file, simply replace the log file name with the `@` character followed by the name of the loghost. For example, to redirect the output normally sent to the `messages`, `secure`, and `maillog` log files, make these changes to the above file:

```
# Log anything (except mail) of level info or higher.
# Don't log private authentication messages!
*.info;mail.none;news.none;authpriv.none;cron.none  @loghost

# The authpriv file has restricted access.
authpriv.*                      @loghost

# Log all the mail messages in one place.
mail.*                          @loghost
```

The messages will now be sent to the `syslogd` running on the computer named loghost. The name loghost was not an arbitrary choice. It is customary to create such a host name and make it an alias to the actual system acting as the loghost. That way, if you ever need to switch the loghost duties to a different machine, you only need to change the loghost alias; you do not need to reedit the `syslog.conf` file on every computer.

Understanding the messages logfile

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Because of the many programs and services that record information to the messages logfile, it is important that you understand the format of this file. Examining this file will often give you a good early warning of problems developing on your system. Each line in the file is a single message recorded by some program or service. Here is a snippet of an actual messages log file:

```
Feb 25 11:04:32 toys network: Bringing up loopback interface:  succeeded
Feb 25 11:04:35 toys network: Bringing up interface eth0:  succeeded
Feb 25 13:01:14 toys vsftpd(pam_unix)[10565]: authentication failure;
    logname= uid=0 euid=0 tty= ruser= rhost=10.0.0.5  user=chris
Feb 25 14:44:24 toys su(pam_unix)[11439]: session opened for
    user root by chris(uid=500)
```

This is really very simple when you know what to look for. Each message is divided into five main parts. From left to right they are:

- The date and time that the message was logged
- The name of the computer that the message came from
- The program or service name that the message pertains to
- The process number (enclosed in square brackets) of the program sending the message
- The actual text message itself

Let's examine the file snippet above. In the first two lines, you can see that I restarted the network. The next line shows that I tried to log in as the user named `chris` to get to the FTP server on this system from a computer at address 10.0.0.5 (I typed the wrong password and authentication failed). The last line shows that I used the `su` command to become root user.

By occasionally reviewing the messages file and the secure file, it is possible to catch a cracking attempt before it is successful. If you see an excessive number of connection attempts for a particular service, especially if they are coming from systems on the Internet, you may be under attack.

Cross-Reference

The logsentry package provides a simpler way to manage your log files and not miss critical messages. Refer to "Monitoring Log Files with LogSentry" later in this chapter.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Monitoring Log Files with LogSentry

Red Hat Linux has the ability to monitor and log nearly every activity that can occur on your computer. On a busy system, massive amounts of informational and error messages are produced and placed in log files. For the administrator, the hard part of monitoring log files isn't detecting or logging security problems; the hard part is remembering to check the log files and sift out those messages that pose a threat from all the other stuff that gets logged. This section describes how to monitor log files with the LogSentry package.

The LogSentry package (formerly called LogCheck) is a handy tool that you can use to easily manage your system log files. Because you're more likely to glance through an e-mail than you are to remember to check log files, LogSentry puts information in front of you that may otherwise go unnoticed. LogSentry checks log files produced by the standard Linux `syslog` facility, attempts to filter out messages that don't represent any security threat, and then categorizes messages that could represent a threat and e-mails those messages to the system administrator. By default, LogSentry will check messages in the `messages`, `secure`, and `mail` log files in the `/var/log` directory. By changing the LogSentry configuration files, you can change which log files are checked, how messages are filtered, and how often log-summary e-mail messages are sent. You may also want to change which features the syslog facility monitors and the level of syslog monitoring if you are interested in debugging a particular system feature.

Note Recently, Red Hat added the logwatch package to the distribution. Logwatch is configured to run once each day to monitor your log files. While logwatch is not as comprehensive a tool as LogSentry, you can expect to see a logwatch e-mail message produced daily that highlights any dangerous system activity that it detects. That e-mail is sent to your computer's system administrator.

Downloading and installing LogSentry

You can find the `logsentry` package at `rpmfind.net` and several FTP sites on the Internet. After `logsentry` is downloaded, to install the package, simply run the following command from the directory you downloaded it to:

rpm -Uhv logsentry*

The LogSentry package consists of four configuration files (in the `/etc/logsentry` directory), the `cron` file that runs `logcheck.sh` (in `/etc/cron.d/logsentry.cron`), and the `logcheck.sh` and `logtail` commands (in `/usr/sbin`). There are also README files in the `/usr/share/doc/logcheck*` directory.

Because there are different versions of LogSentry floating around on the Web, including older versions named logcheck, the locations of commands and instructions for using LogSentry may be different from what I describe here. For example, the original logcheck package may include a `logcheck.sh` script in `/usr/bin` for running logcheck. The logsentry package used here is logsentry-1.1.1-1.i686.rpm.

Setting up LogSentry

You don't have to do anything to get LogSentry working. After LogSentry is installed, it runs once a day at midnight. The results are then e-mailed to the root user on the local host computer. There are several things you can do, however, to tailor LogSentry to suit your particular needs. See the "Configuring LogSentry to suit your needs" section, later in this chapter, for more information.

Running LogSentry

When you install the logsentry package, a script named `logsentry.cron` is placed in the `/etc/cron.d` directory. In this case, the `/etc/cron.d/logsentry.cron` script simply runs the `/usr/sbin/logcheck.sh` script.

Because the `logsentry.cron` script runs hourly, this means that you (or the administrator) will receive 24 e-mail messages each day, each containing the filtered log messages. You can create your own `cron` script to have the `logcheck.sh` launched on any schedule you like.

Note The /usr/sbin/logcheck.sh script uses the /usr/sbin/logtail command to gather only those log messages that you haven't already seen. The logtail command does this by creating a .offset file for each log file that Logcheck monitors. The next time Logcheck is run, only log messages that have arrived since the previous run are checked.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using LogSentry

After LogSentry has been set up and run, to begin using LogSentry you start by simply reading the e-mail that LogSentry sends you. By default, the `root` user on your Red Hat Linux system will receive an e-mail message from LogSentry each hour. Log messages that are matched, and not excluded, are sorted under one of the following three headings in each e-mail message:

Active System Attack Alerts: Represents messages that may represent an attack on your system.

Security Violations: Includes failures and violations that may indicate a problem, but not necessarily an attack on your system.

Unusual System Events: Includes all log messages that are neither matched nor excluded.

The following is an example of a LogSentry e-mail message.

Return-Path: <root@localhost.localdomain>Received: (from root@localhost)

Subject: maple 11/06/03:19.01 ACTIVE SYSTEM ATTACK!

Status: R

Active System Attack Alerts

=-=-=-=-=-=-=-=-=-=-=-=-=-=

Nov 6 18:02:28 maple portsentry[1102]: attackalert: Possible stealth

scan from unknown host to TCP port: 111 (accept failed)

Nov 6 18:33:26 maple sendmail[1863]: f371XJw01863: "wiz" command from

duck.handsonhistory.com [10.0.0.28] (127.0.0.1)

Nov 6 18:33:29 maple sendmail[1863]: f371XJw01863: "debug" command from

duck.handsonhistory.com [10.0.0.28] (127.0.0.1)

Security Violations

=-=-=-=-=-=-=-=-=-=-=

Nov 6 18:02:28 maple portsentry[1102]: attackalert: Possible stealth

scan from unknown host to TCP port: 111 (accept failed)

Nov 6 18:39:14 maple -- root[1121]: ROOT LOGIN ON tty1

Unusual System Events

=-=-=-=-=-=-=-=-=-=-=

Nov 6 18:01:35 maple last message repeated 291877 times

Nov 6 18:14:10 maple last message repeated 297510 times

Nov 6 18:20:58 maple kernel: SB 4.16 detected OK (220)

Nov 6 18:20:58 maple kernel: SB16: Bad or missing 16 bit DMA channel

Nov 6 18:20:58 maple kernel: sb: 1 Soundblaster PnP card(s) found.

Nov 6 18:38:37 maple syslog: syslogd startup succeeded

Nov 6 18:38:37 maple kernel: klogd 1.4-0, log source = /proc/kmsg

started.

Nov 6 18:38:37 maple kernel: Inspecting /boot/System.map-2.4.2-0.1.49

In the preceding e-mail, under the Active System Attack Alerts heading, you can see that a scan of port 111 (portmapper service) was detected by the PortSentry service. The next two messages indicate that a user from the host `duck.handsonhistory.com` tried to scan `sendmail` to see if `debug` and `wiz` services could be accessed. Under the Security Violations heading, the possible stealth scan appeared again. A normal login by the `root` user was also detected (although it represented no particular threat

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

in this case).
Under the Unusual System Events heading, as noted earlier, are included all messages not matched (specifically added to another category) or excluded (specifically ignored) by any of the filter files. A lot of normal system-activity messages appear here. Over time, you may want to explicitly include or exclude messages that you see all the time or that catch your eye as a potential problem you need to watch. For example, the preceding lines that say `last message repeated 291877 times` indicate a potential denial-of-service attack. You may want to add the keywords "message repeated" to your `logcheck.hacking` list. Likewise, the words "Soundblaster PnP card(s) found" could be added to the `logcheck.ignore` file, because that message reflects normal processing.

In general, you want to react to attacks that you detect by preventing an attacker from gaining access to your system. If an attacker does get in, you want to get that attacker out of your system and clean up the damage as best you can. An excellent tool for detecting, logging, and denying access to your system by attackers is the PortSentry package discussed later in this chapter.

Configuring LogSentry to suit your needs

After LogSentry is installed, it will run without requiring any configuration. However, to better suit your needs, there are several configuration files you can modify. The following section describes those files.

Editing the logcheck script

The `/usr/sbin/logcheck.sh` script scans your log files and sorts the log messages that are e-mailed. You can change much of the behavior of the `logcheck.sh` script by changing the values of the variables within the script. To change the behavior of the `logcheck.sh` script, follow these steps:

1. Make a copy of the `/usr/sbin/logcheck.sh` file. For example:
2. # **cp /usr/sbin/logcheck.sh /usr/sbin/logcheck.sh.old**
3. Open the script in any text editor while logged in as the `root` user and make any changes to the script. The following bullet list describes values that you may want to change.

   SYSADMIN: This variable defines the `root` user as the one to receive the e-mail messages resulting from running Logcheck. You can change `root` to anyone you want to receive the LogSentry messages. This can be either local users or users on other computers (that is, *user@hostname*).

   SYSADMIN=root

   TMPDIR: Sets where LogSentry writes its temporary files during processing. The directory is created when LogSentry starts and is removed before Logcheck finishes. You can change the location by modifying the following entry:

   TMPDIR=/tmp/logcheck$$-$RANDOM

   GREP: The logcheck.sh script relies on a `grep` command that supports the `-i`, `-v`, and `-f` options to search the log files. By default, the `egrep` command is used for this purpose. You could change the value to `grep` or another command by modifying the following variable:

   GREP=egrep

   MAIL: E-mail is sent by LogSentry using the `mail` command. To have LogSentry use a command other than the `mail` command to send e-mail messages to the administrator, change the following `mail` variable:

   MAIL=mail

   Filter files: There are four filter files defined by LogSentry. The files each contain keywords that are either used to find messages or exclude messages that contain those keywords. The following variables are set to indicate the location of the four LogSentry filter files:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

HACKING_FILE=/etc/logsentry/logcheck.hacking

VIOLATIONS_FILE=/etc/logsentry/logcheck.violations

VIOLATIONS_IGNORE_FILE=/etc/logsentry/logcheck.violations.ign
ore

IGNORE_FILE=/etc/logsentry/logcheck.ignore
For more information on filter files, see the "Changing LogSentry filter files" section, later in this chapter.

Log files: The last entries in the `/usr/sbin/logcheck.sh` script that you may want to change designate which log files are monitored by the `logcheck.sh` script. By default, the script runs the `logtail` command to check the `messages`, `secure`, and `maillog` files (in `/var/log`). The following lines define which log files are checked and determine where the output of `logtail` is temporarily written.

```
$LOGTAIL /var/log/messages > $TMPDIR/check.$$

$LOGTAIL /var/log/secure >> $TMPDIR/check.$$

$LOGTAIL /var/log/maillog >> $TMPDIR/check.$$
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip If you like, you can have more log files checked by adding more lines like the preceding ones. Just make sure that the first line includes a single arrow (to overwrite a previous ch

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Changing LogSentry filter files

The `/etc/logsentry` directory contains four filter files that define which messages are matched and e-mailed to the administrator. The contents of these files are simply keywords. Log files are searched (or *grepped*) for these keywords and sorted or discarded based on the results of the search.

Some keyword filter files are intended to uncover words or phrases that would appear in a log message in the event of a system break-in or misuse. Other keyword files are intended to find messages that pose no security threat (so that the messages can be excluded from the e-mailed log messages). Messages that match neither the included or excluded keywords are appended to the Unusual System Events heading of the e-mail summary output.

You can use filter files as they are. However, over time, you may want to modify these files for several reasons. If you are receiving repetitive, non-threatening messages in the LogSentry e-mails, you can add keywords that can filter out those messages. Also, you can add keywords later as you learn about new types of security breaches that you may want to look for.

Besides including alphanumeric characters, keywords can also include wildcard characters. For example, you could use an asterisk (*) to match any string of characters, a question mark (?) to match any single character, or a dollar sign ($) to match a keyword that appears at the end of a line.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution Use wildcards carefully. A mistaken wildcard character can result in too many or too few messages being included or excluded.

The four LogSentry filter files in the `/etc/logsentry` directory are:

logcheck.hacking: Contains keywords that appear in log messages that represent known hacking attacks.

logcheck.ignore: Contains keywords that represent messages that should always be ignored.

logcheck.violations: Contains keywords that represent negative activities that may or may not represent real intrusions on your system.

logcheck.violations.ignore: Contains keywords that represent messages that should be ignored from those found as part of the violations check.

Each of these files is described in the following sections.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note It is important to note that messages that are neither explicitly matched (from log check.hacking and log check.violations) nor explicitly excluded (from log check.ignore and log ch

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

logcheck.hacking

Keywords from the `/etc/logsentry/logcheck.hacking` file are meant to uncover log messages representing attacks on your system. Log messages that are matched by keywords in this file are output in e-mail messages to your system administrator under the `logcheck.hacking` heading.

Messages that appear under this heading are the first messages to appear in the e-mail message. You can add other keywords to this file as you learn about messages that represent different types of attacks on your system.
Following are some examples of keywords that appear in the `logcheck.hacking` file:

"wiz"

"WIZ"

"debug"

"DEBUG"

ATTACK

nested

VRFY bbs

VRFY decode

VRFY uudecode

rlogind.*: Connection from .* on illegal port

rshd.*: Connection from .* on illegal port

sendmail.*: user .* attempted to run daemon

uucico.*: refused connect from .*

tftpd.*: refused connect from .*

login.*: .*LOGIN FAILURE.* FROM .*root

login.*: .*LOGIN FAILURE.* FROM .*guest

Most of the messages in this file are used to match log messages that result from someone probing your system with the Internet Security Scanner (ISS). ISS is a tool that can scan a set of IP addresses for potential security weaknesses. Though most of the security holes ISS checks for have been plugged over time, these log messages alert you to the fact that someone is checking the security of your system.
The `debug` and `wiz` keywords shown in the above file will catch attempts by ISS to access `wiz` and `debug` services from the `sendmail` service. Likewise, `VRFY` is a command that ISS sends to requests of the `sendmail` service to ask for different user names. Other keyword phrases shown in the previous example match failed attempts to connect using different Linux network services (such as `rlogind`, `rshd`, `sendmail`, `uucico`, and so on).

logcheck.ignore

Keywords in the `/etc/logsentry/logcheck.ignore` file are used to find log messages that should be excluded (that is, ignored) by LogSentry and will therefore not appear in e-mail summaries. The keywords in this file reduce the number of log messages that are e-mailed to the administrator, making it easier to find the real problems. The following are some examples of keywords from the `logcheck.ignore` file.

cron.*CMD

cron.*RELOAD

cron.*STARTUP

ftp-gw.*: exit host

ftp-gw.*: permit host

ftpd.*ANONYMOUS FTP LOGIN

http-gw.*: exit host

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

http-gw.*: permit host

identd.*Successful lookup

identd.*from:

named.*Response from
Most of the entries in this file are used to match messages that represent the normal operation of various system services. The keywords shown in the previous example represent normal processing of `cron`, `ftp`, `http`, `identd`, and `named` features.

Notice that all of the keyword phrases shown in the preceding example include an asterisk (*) wildcard. If you add your own keywords to this file, using asterisks and other wildcards can help you be specific about the log messages you exclude.

logcheck.violations

There are certain words that imply negative behavior occurring on your computer. Though these words may not be associated with any particular attack, LogSentry notes messages containing these words and displays them under the following heading and e-mails them to the system administrator:

Security Violations

=-=-=-=-=-=-=-=-=-=-
The following is an example of some of the keywords that appear in the `logcheck.violations` file:

ATTACK

BAD

DEBUG

FAILURE

ILLEGAL

REFUSED

denied

failed

unapproved

attackalert

Adding your own keywords to the file can help flag log messages that may be of particular concern for your computer. For example, you can add keywords that represent improper use of services that are not standard Linux features but can be accessed from the network.

logcheck.violations.ignore

Use the `logcheck.violations.ignore` file to exclude messages that were matched from the `logcheck.violations` file, but are known to not represent security problems. By default, only the following entry is contained in this file:

stat=Deferred
The previous keyword causes LogSentry to ignore log messages from `sendmail` that represent e-mail messages that haven't been sent because the receiving server was temporarily unavailable. As you use LogSentry, you are likely to repeatedly encounter certain log messages that represent no security threat. Add keywords here (specifically as possible) to exclude those messages from appearing continuously in e-mail messages from LogSentry.

Modifying syslog

The `syslog` service gathers the log messages of system activity that are used by LogSentry. Red Hat Linux, as well as most other Linux and UNIX systems, comes with `syslog` installed and operational by default. You can modify the `/etc/syslog.conf` file to tailor the behavior of `syslog` to best suit the way you use your system.
The `syslog` service is part of the `sysklogd` software package. To make sure that `sysklogd` is installed, type the following at a shell prompt:

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

rpm -q sysklogd

The `syslogd` service is started automatically from the `/etc/init.d/syslog` start-up script. After that script has been run on your system, two daemon processes should be active on your system: `syslogd` and `klogd`. To see if they are running, type the following:

ps -ax | grep log

The `/etc/syslog.conf` file contains information that defines which activities are logged, and from which system. LogSentry monitors three of the log files created by `syslog` (in the `/var/log` directory): `messages`, `secure`, and `maillog`. The following lines in the `/etc/syslog.conf` file instruct `syslog` to create those log files:

```
*.info;mail.none;news.none;authpriv.none;cron.none   /var/log/messages
authpriv.*                                           /var/log/secure
mail.*                                               /var/log/maillog
```

Services on your Red Hat Linux system produce messages of different levels. Message levels, from most critical to least critical, are listed in Table 14-7.

Table 14-7: Message Levels

Level	What it Means
alert	immediate action needed
crit	Critical
debug	detailed processing information
emerg	system unusable
err	error condition
info	purely informational
notice	important, but not an error
warning	potential error

The line shown in the example indicates that all messages from the info level (`*.info`) and above are logged to the `/var/log/messages` file. However, messages of types `mail`, `news`, `authpriv`, and `cron` are excluded because they are sent to other log files. All `authpriv` (`authpriv.*`) messages are logged to the `/var/log/secure` file. `mail` messages (`mail.*`) are all logged to the `/var/log/maillog` file.

With this default configuration of `syslog`, LogSentry should catch all major security related activities. There are a few situations, however, where you may want to modify the `/etc/syslog.conf` file. For example, if you are receiving a lot of log messages for a particular type of service (such as `ppp` if you are having trouble with a dial-up connection), you may consider directing messages for that service to its own log file. Then, if LogSentry uncovers a problem, it's easier to go through only that log file for those messages relating to the problem service.

Another temporary change you may consider is if you need to debug a problem with your system. Changing `*.info` to `*.debug` temporarily can give you more details on a problem. (Make sure to change it back later, or `syslog` will chew up too much system resources.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using Password Protection

Passwords are the most fundamental security tool of any modern operating system and consequently, the most commonly attacked security feature. It is natural to want to choose a password that is easy to remember, but very often this means choosing a password that is also easy to guess. Crackers know that on any system with more than a few users, at least one person is likely to have an easily guessed password.

By using the "brute force" method of attempting to log in to every account on the system and trying the most common passwords on each of these accounts, a persistent cracker has a good shot of finding a way in. Remember that a cracker will automate this attack, so thousands of login attempts are not out of the question. Obviously, choosing good passwords is the first and most important step to having a secure system.

Here are some things to avoid when choosing a password:

- Do not use any variation of your login name or your full name. Even if you use varied case, append or prepend numbers or punctuation, or type it backwards, this will still be an easily guessed password.
- Do not use a dictionary word, even if you add numbers or punctuation to it.
- Do not use proper names of any kind.
- Do not use any contiguous line of letters or numbers on the keyboard (such as "qwerty" or "asdfg").

Choosing good passwords

A good way to choose a strong password is to take the first letter from each word of an easily remembered sentence. The password can be made even better by adding numbers, punctuation, and varied case. The sentence you choose should have meaning only to you, and should not be publicly available (choosing a sentence on your personal Web page is a bad idea). Table 14-8 lists examples of strong passwords and the tricks used to remember them.

Table 14-8: Ideas for Good Passwords

Password	How to Remember it
`Mrci7yo!`	My rusty car is 7 years old!
`2emBp1ib`	2 elephants make BAD pets, 1 is better
`ItMc?Gib`	Is that MY coat? Give it back

The passwords look like gibberish, but are actually rather easy to remember. As you can see, I can place emphasis on words that stand for capital letters in the password. You set your password using the `passwd` command. Type the `passwd` command within a command shell, and it will enable you to change your password. First, it will prompt you to enter your old password. To protect against someone "shoulder surfing" and learning your password, the password will not be displayed as you type. Assuming you type your old password correctly, the `passwd` command will prompt you for the new password. Once you type in your new password, the `passwd` command checks the password against `cracklib` to determine if it is a *good* or *bad* password. Non-root users will be required to try a different password if the one they have chosen is not a good password. The root user is the only user who is permitted to assign *bad* passwords. Once the password has been accepted by `cracklib`, the `passwd` command will ask you to enter the new password a second time to make sure there are no typos (which are hard to detect when you can't see what you are typing). When running as root, it is possible to change a user's password by supplying that user's login name as a parameter to the `passwd` command. Typing this:

passwd joe

results in the `passwd` command prompting you for joe's new password. It does not prompt you for his old password in this case. This allows root to reset a user's password when that user has forgotten it (an event that happens all too often).

Using a shadow password file

In early versions of UNIX, all user account and password information was stored in a file that all users

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

could read (although only root could write to it). This was generally not a problem because the password information was encrypted. The password was encrypted using a *trapdoor algorithm*, meaning the non-encoded password could be encoded into a scrambled string of characters, but the string could not be translated back to the non-encoded password.

How does the system check your password in this case? When you log in, the system encodes the password you entered, compares the resulting scrambled string with the scrambled string that is stored in the password file, and grants you access only if the two match. Have you ever asked a system administrator what the password on your account is only to hear, "I don't know" in response? If so, this is why: The administrator really doesn't have the password, only the encrypted version. The non-encoded password exists only at the moment you type it.

Breaking encrypted passwords

There is a problem with people being able to see encrypted passwords, however. Although it may be difficult (or even impossible) to reverse the encryption of a trapdoor algorithm, it is very easy to encode a large number of password guesses and compare them to the encoded passwords in the password file. This is, in orders of magnitude, more efficient than trying actual login attempts for each user name and password. If a cracker can get a copy of your password file, the cracker has a much better chance of breaking into your system.
Fortunately, Linux and all modern UNIX systems support a shadow password file by default. The shadow file is a special version of the `passwd` file that only root can read. It contains the encrypted password information, so passwords can be left out of the `passwd` file, which any user on the system can read. Linux supports both the older, single password file method as well as the newer shadow password file. You should always use the shadow password file (it is used by default).

Checking for the shadow password file

The password file is named `passwd` and can be found in the `/etc` directory. The shadow password file is named shadow and is also located in `/etc`. If your `/etc/shadow` file is missing, then it is likely that your Linux system is storing the password information in the `/etc/passwd` file instead. Verify this by displaying the file with the `less` command.

less /etc/passwd

Something similar to the following should be displayed:

root:DkkS6Uke799fQ:0:0:root:/root:/bin/bash

bin:*:1:1:bin:/bin:

daemon:*:2:2:daemon:/sbin:

.

.

.

mary:KpRUp2ozmY5TA:500:100:Mary Smith:/home/mary:/bin/sh

joe:0sXrzvKnQaksI:501:100:Joe Johnson:/home/joe:/bin/sh

jane:ptNoiueYEjwX.:502:100:Jane Anderson:/home/jane:/bin/sh

bob:Ju2vY7A0X6Kzw:503:100:Bob Renolds:/home/bob:/bin/sh
Each line in this listing corresponds to a single user account on the Linux system. Each line is made up of seven fields separated by colon (:) characters. From left to right the fields are the login name, the encrypted password, the user ID, the group ID, the description, the home directory, and the default shell. Looking at the first line, you see that it is for the root account and has an encrypted password of DkkS6Uke799fQ. We can also see that root has a user ID of zero, a group ID of zero, and a home directory of `/root`, and root's default shell is `/bin/sh`.
All of these values are quite normal for a root account, but seeing that encrypted password should set off alarm bells in your head. It confirms that your system is not using the shadow password file. At this point, you should immediately convert your password file so that it uses `/etc/shadow` to store the password information. You do this by using the `pwconv` command. Simply log in as root (or use the `su` command to become root) and enter the `pwconv` command at a prompt. It will print no messages, but when your shell prompt returns, you should have a `/etc/shadow` file and your `/etc/passwd` file

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

should now look like this:

root:x:0:0:root:/root:/bin/bash

bin:x:1:1:bin:/bin:

daemon:x:2:2:daemon:/sbin:

.

.

.

mary:x:500:100:Mary Smith:/home/mary:/bin/sh

joe:x:501:100:Joe Johnson:/home/joe:/bin/sh

jane:x:502:100:Jane Anderson:/home/jane:/bin/sh

bob:x:503:100:Bob Renolds:/home/bob:/bin/sh

Encrypted password data is replaced with an x. Password data moved to `/etc/shadow`.

There is also a screen-oriented command called `authconfig` that you can use to manage shadow passwords and other system authentication information. This tool also has features that let you work with MD5 passwords, LDAP authentication, or Kerberos 5 authentication as well. Type `authconfig` and step through the screens to use it.

To work with passwords for groups, you can use the `grpconv` command to convert passwords in `/etc/groups` to shadowed group passwords in `/etc/gshadow`. If you change `passwd` or group passwords and something breaks (you are unable to log in to the accounts), you can use the `pwunconv` and `grpunconv` commands, respectively, to reverse password conversion.

So, now you are using the shadow password file and picking good passwords. You have made a great start toward securing your system. You may also have noticed by now that security is not just a one-time job. It is an ongoing process, as much about policies as programs. Keep reading to learn more.

Team LiB

PREVIOUS NEXT

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using Encryption Techniques

The previous sections told you how to lock the doors to your Red Hat Linux system to deny access to crackers. The best dead bolt lock, however, is useless if you are mugged in your own driveway and have your keys stolen. Likewise, the best computer security can be for naught if you are sending passwords and other critical data unprotected across the Internet.

A savvy cracker can use a tool called a protocol analyzer or a network sniffer to peek at the data flowing across a network and pick out passwords, credit card data, and other juicy bits of information. The cracker does this by breaking into a poorly protected system on the same network and running software, or by gaining physical access to the same network and plugging in his or her own equipment.

You can combat this sort of theft by using encryption. The two main types of encryption in use today are symmetric cryptography and public-key cryptography.

Symmetric cryptography

Symmetric cryptography, also called private-key cryptography, uses a single key to both encrypt and decrypt a message. This method is generally inappropriate for securing data that is expected to be utilized by a third party, due to the complexity of secure key exchange. Symmetric cryptography is generally useful for encrypting data for one's own purposes.

A classic use of symmetric cryptography is for a personal password vault. Anyone who has been using the Internet for any amount of time has accumulated a quantity of user names and passwords for accessing various sites and resources. A personal password vault lets you store this access information in an encrypted form. The end result is that you only have to remember one password to unlock all of your access information.
Until recently, the United States government was standardized on a symmetric encryption algorithm called DES (Data Encryption Standard) to secure important information. Because there is no direct way to crack DES encrypted data, to decrypt DES encrypted data without a password you would have to use an unimaginable amount of computing power to try to guess the password. This is also known as the *brute force* method of decryption.

As personal computing power has increased nearly exponentially, the DES algorithm has had to be retired. In its place, after a very long and interesting search, the United States. government has accepted the Rijndael algorithm as what it calls the AES (Advanced Encryption Standard). Although the AES algorithm is also subject to brute force attacks, it requires significantly more computing power to crack than the DES algorithm does.
For more information on AES, including a command line implementation of the algorithm, you can visit `http://aescrypt.sourceforge.net/.`

Public-key cryptography

Public-key cryptography does not suffer from key distribution problems, and that is why it is the preferred encryption method for secure Internet communication. This method uses two keys, one to encrypt the message and another to decrypt the message. The key used to encrypt the message is called the public key because it is made available for all to see. The key used to decrypt the message is the private key and is kept hidden. The entire process works like this:

Imagine that you want to send me a secure message using public-key encryption. Here is what we need:

1. I must have a public and private key pair. Depending on the circumstances, I may generate the keys myself (using special software) or obtain the keys from a key authority.
2. You wish to send me a message, so you first look up my public key (or more accurately, the software you are using looks it up).
3. You encrypt the message with the public key. At this point, the message can only be decrypted with the private key (the public key cannot be used to decrypt the message).
4. I receive the message and use my private key to decrypt it.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Secure Socket Layer

A classic implementation of public-key cryptography is with SSL (secure socket layer) communication. This is the technology that enables you to securely submit your credit card information to an online merchant. The elements of an SSL encrypted session are as follows:

SSL Enabled Web Browser (Mozilla, Internet Explorer, Opera, Konquerer, etc.)
SSL Enabled Web Server (Apache)
SSL Certificate

To initiate an SSL session, a Web browser first makes contact with a Web server on port 443, also known as the HTTPS port (Hypertext Transport Protocol Secure). Once a socket connection has been established between the two machines, the following occurs:

1. Server sends SSL certificate to browser.
2. Browser verifies identity of server through SSL certificate.
3. Browser generates symmetric encryption key.
4. Browser uses SSL certificate to encrypt symmetric encryption key.
5. Browser sends encrypted key to the server.
6. Server decrypts the symmetric key with its private key counterpart of the public SSL certificate.
7. Browser and server can now encrypt and decrypt traffic based on a common knowledge of the symmetric key.

Secure data interchange can now occur.

Creating SSL Certificates

In order to create your own SSL certificate for secure HTTP data interchange, you must first have an SSL-capable Web server. The Apache Web server (httpd package), which comes with Red Hat Linux, is SSL-capable. Once you have a server ready to go, you should familiarize yourself with the important server-side components of an SSL certificate:

ls -l /etc/httpd/conf

drwxr-xr-x 7 root root 4096 Aug 12 23:45 .

drwxr-xr-x 4 root root 4096 Aug 13 00:23 ..

-rw-r--r-- 1 root root 35918 Jul 14 15:45 httpd.conf

lrwxrwxrwx 1 root root 37 Aug 12 23:45 Makefile ->

../../../usr/share/ssl/certs/Makefile

drwx------ 2 root root 4096 Aug 12 23:45 ssl.crl

drwx------ 2 root root 4096 Aug 12 23:45 ssl.crt

drwx------ 2 root root 4096 Jul 14 15:45 ssl.csr

drwx------ 2 root root 4096 Aug 12 23:45 ssl.key

drwx------ 2 root root 4096 Jul 14 15:45 ssl.prm

ls -l /etc/httpd/conf.d/ssl.conf

-rw-r--r-- 1 root root 11140 Jul 14 15:45 ssl.conf

The `/etc/httpd/conf` and `/etc/httpd/conf.d` directories contains all of the components necessary to create your SSL certificate. Each component is defined as follows:

httpd.conf — Web server configuration file.
Makefile — Certificate building script.
ssl.crl — Certificate revocation list directory.
ssl.crt — SSL certificate directory.
ssl.csr — Certificate service request directory.
ssl.key — SSL certificate private key directory.
ssl.prm — SSL certificate parameters.
ssl.conf — Primary Web server SSL configuration file.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Now that you're familiar with the basic components, take a look at the tools used to create SSL certificates:

cd /etc/httpd/conf

make

This makefile allows you to create:

- o public/private key pairs
- o SSL certificate signing requests (CSRs)
- o self-signed SSL test certificates

To create a key pair, run "make SOMETHING.key".

To create a CSR, run "make SOMETHING.csr".

To create a test certificate, run "make SOMETHING.crt".

To create a key and a test certificate in one file, run "make SOMETHING.pem".

To create a key for use with Apache, run "make genkey".

To create a CSR for use with Apache, run "make certreq".

To create a test certificate for use with Apache, run "make testcert".

Examples:

- make server.key
- make server.csr
- make server.crt
- make stunnel.pem
- make genkey
- make certreq
- make testcert

The `make` command utilizes the Makefile to create SSL certificates. Without any arguments the `make` command simply prints the information listed above. The following defines each argument you can give to `make`:

make server.key — Creates generic public/private key pairs.
make server.csr — Generates a generic SSL certificate service request.
make server.crt — Generates a generic SSL test certificate.
make stunnel.pem — Generates a generic SSL test certificate, but puts the private key in the same file as the SSL test certificate.
make genkey — Same as `make server.key` except it places the key in the `ssl.key` directory.
make certreq — Same as `make server.csr` except it places the certificate service request in the `ssl.csr` directory.
make testcert — Same as `make server.crt` except it places the test certificate in the `ssl.crt` directory.

Using third-party certificate signers

In the real world, I know who you are because I recognize your face, your voice and your mannerisms. On the Internet, I cannot see these things and must rely on a trusted third party to vouch for your identity. To ensure that a certificate is immutable, it has to be signed by a trusted third party when the certificate is issued and validated every time an end user taking advantage of your secure site loads it. The following is a list of the trusted third-party certificate signers:

GlobalSign — https://www.globalsign.net/

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Baltimore — https://www.baltimore.com/
GeoTrust — https://www.geotrust.com/
VeriSign — https://www.verisign.com/
FreeSSL — http://www.freessl.com/
Thawte — http://www.thawte.com/
EnTrust — http://www.entrust.com/
ipsCA — http://www.ipsca.com/
COMODO Group — http://www.comodogroup.com/

Each of these certificate authorities has gotten a chunk of cryptographic code embedded into nearly every Web browser in the world. This chunk of cryptographic code allows a Web browser to determine whether or not an SSL certificate is authentic. Without this validation, it would be trivial for crackers to generate their own certificates and dupe people into thinking they are giving sensitive information to a reputable source.

Certificates that are not validated are called *self-signed certificates*. If you come across a site that has not had its identity authenticated by a trusted third party, your Web browser will display a message similar to the one shown in Figure 14-4.

Figure 14-4: A pop-up window alerts you when a site is not authenticated.

This does not necessarily mean that you are encountering anything illegal, immoral or fattening. Many sites opt to go with *self-signed* certificates, not because they are trying to pull a fast one on you, but because there may not be any reason to validate the true owner of the certificate and they do not wish to pay the cost of getting a certificate validated. Some reasons for a using a *self-signed* certificate include:

The Web site accepts no input. In this case, you as the end user, have nothing to worry about. There is no one trying to steal your information because you aren't giving out any information. Most of the time this is done simply to secure the Web transmission from the server to you. The data in and of itself may not be sensitive, but, being a good netizen, the site has enabled you to secure the transmission to keep third parties from sniffing the traffic.

The Web site caters to a small clientele. If you run a Web site that has a very limited set of customers, such as an Application Service Provider, you can simply inform your users that you have no certificate signer and that they can browse the certificate information and validate it with you over the phone or in person.

Testing. It makes no sense to pay for an SSL certificate if you are only testing a new Web site or Web-based application. Use a *self-signed* certificate until you are ready to go live.

Creating a Certificate Service Request

To create a third party validated SSL certificate, you must first start with a CSR (Certificate Service Request). To create a CSR, do the following on your Web server:

cd /etc/httpd/conf

make certreq

umask 77 ; \

/usr/bin/openssl genrsa -des3 1024 > /etc/httpd/conf/ssl.key/server.key

.

.

.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

You will now be asked to enter a password to secure your private key. This password should be at least eight characters long, and should not be a dictionary word or contain numbers or punctuation. The characters you type will not appear on the screen, in order to prevent someone from shoulder surfing your password.

Enter pass phrase:

Enter the password once again to verify.

Verifying - Enter pass phrase:

The certificate generation process now begins.

At this point, it is time to start adding some identifying information to the certificate that the third party source will later validate. Before you can do this, you must unlock the private key you just created. Do so by typing the password you typed above. Then enter information as you are prompted. An example of a session for adding information for your certificate is shown below:

Enter pass phrase for /etc/httpd/conf/ssl.key/server.key:

You are about to be asked to enter information that will be incorporated

into your certificate request.

What you are about to enter is what is called

a Distinguished Name or a DN.

There are quite a few fields but you can leave some blank

For some fields there will be a default value,

If you enter '.', the field will be left blank.

Country Name (2 letter code) [GB]:**US**

State or Province Name (full name) [Berkshire]: **Connecticut**

Locality Name (eg, city) [Newbury]: **Mystic**

Organization Name (eg, company) [My Company Ltd]:**Acme Marina, Inc.**

Organizational Unit Name (eg, section) []:**InfoTech**

Common Name (eg, your name or your server's hostname)

[]**:www.acmemarina.com**

Email Address []: **webmaster@acmemarina.com**

To complete the process, you will be asked if you want to add any extra attributes to your certificate. Unless you have a reason to provide more information, you should simply hit enter at each of the following prompts to leave them blank.

Please enter the following 'extra' attributes

to be sent with your certificate request

A challenge password []:

An optional company name []:

Getting the CSR Signed

Once your CSR has been created, you now need to send it to a signing authority for validation. The first step in this process is to select a signing authority. Each signing authority has different deals, prices and products. Check out each of the signing authorities listed in the "Using third-party certificate signers " section above to determine which works best for you. The following are areas where signing authorities differ:

- Credibility and stability
- Pricing
- Browser recognition
- Warranties

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Support
Certificate strength

For good comparisons, studies and inside information to make the job of finding an SSL signer easier, go to http://www.whichssl.org.

Once you have selected your certificate signer, you will have to go through some validation steps. Each signer has a different method of validating identity and certificate information. Some require that you fax articles of incorporation, while others require a company officer be made available to talk to a validation operator. At some point in the process you will be asked to copy and paste the contents of the CSR you created into the signer's Web form.

cd /etc/httpd/conf/ssl.csr

cat server.csr

-----BEGIN CERTIFICATE REQUEST-----
MIIB6jCCAVMCAQAwgakxCzAJBgNVBAYTAlVTMRQwEgYDVQQIEwtDb25uZWN0aWN1
dDEPMA0GA1UEBxMGTXlzdGljMRowGAYDVQQKExFBY21lIE1hcmluYSwgSW5jLjER
MA8GA1UECxMISW5mb1RlY2gxGzAZBgNVBAMTEnd3dy5hY21lbWFyaW5hLmNvbTEn
MCUGCSqGSIb3DQEJARYYd2VibWFzdGVyQGFjbWVtYXJpbmEuY29tMIGfMA0GCSqG
SIb3DQEBAQUAA4GNADCBiQKBgQDcYH4pjMxKMldyXRmcoz8uBVOvwlNZHyRWw8ZG
u2eCbvgi6w4wXuHwaDuxbuDBmw//Y9DMI2MXg4wDq4xmPi35EsO1Ofw4ytZJn1yW
aU6cJVQro46OnXyaqXZOPiRCxUSnGRU+0nsqKGjf7LPpXv29S3QvMIBTYWzCkNnc
gWBwwwIDAQABoAAwDQYJKoZIhvcNAQEEBQADgYEANv6eJOaJZGzopNR5h2YkR9Wg
l8oBl3mgoPH60Sccw3pWsoW4qbOWq7on8dS/++QOCZWZI1gefgaSQMInKZ1II7Fs
YIwYBgpoPTMC4bp0ZZtURCyQWrKIDXQBXw7BlU/3A25nvkRY7vgNL9Nq+7681EJ8
W9AJ3PX4vb2+ynttcBI=
-----END CERTIFICATE REQUEST-----

You can use your mouse to copy and paste the CSR into the signer's Web form.

Once you have completed the information validation, paid for the signing and answered all of the questions, you have completed most of the process. Within 48 to 72 hours you should receive an e-mail with your shiny new SSL certificate in it. The certificate will look similar to the following:

-----BEGIN CERTIFICATE-----
MIIEFjCCA3+gAwIBAgIQMI262Zd6njZgN97tJAVFODANBgkqhkiG9w0BAQQFADCB
ujEfMB0GA1UEChMWVmVyaVNpZ24gVHJ1c3QgTmV0d29yazEXMBUGA1UECxMOVmVy
aVNpZ24sIEluXy4xMzAxBgNVBAsTKlZlcmlTaWduIEludGVybmF0aW9uYWwgU2Vy
dmVyIENBIC0gZ2xhc3MgMzFJMEcG10rY2g0Dd3d3LnZlcmlzaWduLmNvbS9DUFMg
SW5jb3JwLmJ51FJlZi4gTElBQklMSVRZIExURC4oYyk5NyBWZXJpU2lnbjAeFw0w
MzAxMTUwMDAwMDBaFw0wNDAxMTUyMzU5NTlaMIGuMQswCQYDVQQGEwJVUzETMBEG
A1UECBMKV2FzaG1uZ3RvHiThErE371UEBxQLRmVkZXJhbCBXYXkxGzAZBgNVBAoU
EklETSBTZXJ2aWMlcywgSW5jLjEMMAoGA1UECxQDd3d3MTMwMQYDVQQLFCpUZXJt
cyBvZiB1c2UgYXQgd3d3LnZlcmlzawduLmNvbS9ycGEgKGMpMDAxFDASBgNVBAMU
C2lkbXNlcnYuY29tMIGfMA0GCSqGS1b3DQEBAQUAA4GNADCBiQKBgQDaHSk+uzOf
7jjDFEnqT8UBa1L3yFILXFjhj3XpMXLGWzLmkDmdJjXsa4x7AhEpr1ubuVNhJVI0
FnLDopsx4pyr4n+P8FyS4M5grbcQzy2YnkM2jyqVF/7yOW2pDl30t4eacYYaz4Qg
q9pTxhUzjEG4twvKCAFWfuhEoGu1CMV2qQ1DAQABo4IBJTCCASEwCQYDVR0TBAIw
ADBEBgNVHSAEPTA7MDkGC2CGSAGG+EUBBxcDMCOwKAYIKwYBBQUHAgEWHGh0dHBz
Oi8vd3d3LnZlcmlzaWduLmNvbS9ycGEwCwYDVRRPBAQDAgWgMCgGA1UdJQQhMB8G
CWCGSAGG+EIEM00c0wIYBQUHAwEGCCsGAQUFBwMCMDQGCCsGAQUFBwEBBCgwJjAk

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

BggrBgEFBQcwAYYYaHR0cDovL29jc2AudmVyaXNpZ24uY29tMEYGA1UdHwQ/MD0w

O6A5oDeGNWh0dHA6Ly9jcmwudmVyaxNpZ24uY29tL0NsYXNzM0ludGVybmF0aW9u

YWxTZXJ2ZXIuY3JsMBkGCmCGSAgG+E+f4Nfc3zYJODA5NzMwMTEyMA0GCSqGSIb3

DQEBBAUAA4GBAJ/PsVttmlDkQai5nLeudLceb1F4isXP17B68wXLkIeRu4Novu13

8lLZXnaR+acHuCkW01b3rQPjgv2y1mwjkPmC1WjoeYfdxH7+Mbg/6fomnK9auWAT

WF0iFW/+a8OWRYQJLMA2VQOVhX4znjpGcVNY9AQSHm1UiESJy7vtd1iX

-----END CERTIFICATE-----

Copy and paste this certificate into an empty file called `server.crt`, which must reside in the `/etc/httpd/conf/ssl.crt` directory, and restart your Web server:

service httpd restart

Assuming your Web site was previously working fine, you can now view it in a secure fashion by placing an "**s**" after the **http** in the Web address. So if you previously viewed your Web site at `http://www.acmemarina.com`, you can now view it in a secure fashion by going to `https://www.acmemarina.com`.

Creating Self-Signed Certificates

Generating and running a self-signed SSL certificate is much easier than having a signed certificate. To generate a self-signed SSL certificate, do the following:

1. Remove the key and certificate that currently exist:
2. # cd /etc/httpd/conf
3. # **rm ssl.key/server.key ssl.crt/server.crt**
4. Create your own server key:
5. # **/usr/bin/openssl genrsa 1024 > ssl.key/server.key**
6. Make the server.key file readable and writable only by root:
7. # **chmod 600 ssl.key/server.key**
8. Create the self-signed certificate by typing the following:
9. # **make testcert**
10. umask 77 ; \
11. /usr/bin/openssl req -new -key /etc/httpd/conf/ssl.key/server.key
12. -x509 -days 365 -out /etc/httpd/conf/ssl.key/server.crt
13. .
14. .
    .

At this point, it is time to start adding some identifying information to the certificate that the third party source will later validate. Before you can do this, you must unlock the private key you just created. Do so by typing the password you typed above. Then follow the sample procedure that follows:

You are about to be asked to enter information that will be

incorporated into your certificate request.

What you are about to enter is what is called

a Distinguished Name or a DN.

There are quite a few fields but you can leave some blank

For some fields there will be a default value,

If you enter '.', the field will be left blank.

Country Name (2 letter code) [GB]:**US**

State or Province Name (full name) [Berkshire]: **Ohio**

Locality Name (eg, city) [Newbury]: **Cincinnati**

Organization Name (eg, company) [My Company Ltd]:**Industrial Press, Inc.**

Organizational Unit Name (eg, section) []:**IT**

Common Name (eg, your name or your server's hostname)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

[]:**www.industrialpressinc.com**

Email Address []: **webmaster@industrialpressinc.com**
The generation process above places all files in the proper place. All you need to do is restart your Web server and add `https` instead of `http` in front of your URL. Don't forget, you'll get a certificate validation message from your Web browser, which you can safely ignore.

Restarting your Web server

By now you've probably noticed that your Web server requires you to enter your certificate password every time it is started. This is to prevent someone from breaking into your server and stealing your private key. Should this happen, you are safe in the knowledge that the private key is a jumbled mess. The cracker will not be able to make use of it. Without such protection, a cracker could get your private key and easily masquerade as you, appearing to be legitimate in all cases.

If you just cannot stand having to enter a password every time your Web server starts, and are willing to accept the increased risk, you can remove the password encryption on your private key. Simply do the following:

cd /etc/httpd/conf/ssl.key

/usr/bin/openssl rsa -in server.key -out server.key

Troubleshooting your certificates

The following tips should help if you are having problems with your SSL certificate.

- Only one SSL certificate per IP address is allowed. If you want to add more than one SSL enabled Web site to your server, you must bind another IP address to the network interface.
- Make sure the permission mask on the `/etc/httpd/conf/ssl.*` directories and their contents is 700 (`rwx------`).
- Make sure you aren't blocking port 443 on your Web server. All `https` requests come in on port 443. If you are blocking it, you will not be able to get secure pages.
- The certificate only lasts for one year. When that year is up, you have to renew your certificate with your certificate authority. Each certificate authority has a different procedure for doing this; check the authority's Web site for more details.
- Make sure you have the `mod_ssl` package installed. If it is not installed, you will not be able to serve any SSL enabled traffic.

Exporting encryption technology

Before describing how to use the various encryption tools, I need to warn you about an unusual policy of the United States government. For many years, the United States government treated encryption technology like munitions. As a result, anyone wishing to export encryption technology had to get an export license from the Commerce Department. This applied not only to encryption software developed within the United States, but also to software obtained from other countries and then re-exported to another country (or even to the same country you got it from). Thus, if you installed encryption technology on your Linux system and then transported it out of the country, you were violating federal law! Furthermore, if you e-mailed encryption software to a friend in another country or let him or her download it from your server, you violated the law.

In January 2000, U.S. export laws relating to encryption software were relaxed considerably. However, often the U.S. Commerce Department's Bureau of Export Administration requires a review of encryption products before they can be exported. U.S. companies are also still not allowed to export encryption technology to countries classified as supporting terrorism.

Using the Secure Shell package

The Secure Shell package (SSH) is a package that provides shell services similar to the `rsh`, `rcp`, and `rlogin` commands, but encrypts the network traffic. It uses Private-Key Cryptography, so it is ideal for use with Internet connected computers. The Red Hat Linux distribution contains the following client and server software packages for SSH: openssh, openssh-client, and openssh-server packages.

Starting the SSH service

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

If you have installed the openssh-server software package, the SSH server is automatically configured to start. The SSH daemon is started from the `/etc/init.d/sshd` start-up script. To make sure the service is set up to start automatically, type the following (as root user):

chkconfig --list sshd

sshd 0:off 1:off 2:on 3:on 4:on 5:on 6:off

This shows that the `sshd` service is set to run in system states 2, 3, 4, and 5 (normal bootup states) and set to be off in all other states. You can turn on the SSH service, if it is off, for your default run state, by typing the following as root user:

chkconfig sshd on

This line turns on the SSH service when you enter run levels 2, 3, 4, or 5. To start the service immediately, type the following:

/etc/init.d/sshd start

Using the ssh, sftp, and scp commands

Three commands you can use with the SSH service are `ssh`, `sftp`, and `scp`. Remote users use the `ssh` command to login to your system securely. The `scp` command lets remote users copy files to and from a system. The `sftp` command provides a safe way to access FTP sites.

Like the normal remote shell services, secure shell looks in the `/etc/hosts.equiv` file and in a user's `.rhost` file to determine whether it should allow a connection. It also looks in the ssh-specific files `/etc/shosts.equiv` and `.shosts`. Using the `shosts.equiv` and the `.shosts` files is preferable because it avoids granting access to the nonencrypted remote shell services. The `/etc/shosts.equiv` and `.shosts` files are functionally equivalent to the traditional `hosts.equiv` and `.rhosts`, so the same instructions and rules apply.

Now you are ready to test the SSH service. From another computer on which SSH has been installed (or even from the same computer if another is not available), type the `ssh` command followed by a space and the name of the system you are connecting to. For example, to connect to the system ratbert.glaci.com, type:

ssh ratbert.glaci.com

If this is the first time you have ever logged in to that system using `ssh`, it will ask you to confirm that you really want to connect. Type **yes** and press Enter when it asks this:

Host key not found from the list of known hosts.

Are you sure you want to continue connecting (yes/no)?

It should then prompt you for a user name and password in the normal way. The connection will then function like a normal telnet connection. The only difference is that the information is encrypted as it travels over the network. You should now also be able to use the `ssh` command to run remote commands.

The `scp` command is similar to the `rcp` command for copying files to and from Linux systems. Here is an example of using the `scp` command to copy a file called `memo` from the home directory of the user named jake to the `/tmp` directory on a computer called maple:

$ **scp /home/jake/memo maple:/tmp**

passwd: **********

memo 100%|**********| 153 0:00

Enter the password for your user name (if a password is requested). If the password is accepted, the remote system indicates that the file has been copied successfully.

Similarly, the `sftp` command starts an interactive FTP session with an FTP server that supports SSH connections. Many security-conscious people prefer `sftp` to other `ftp` clients because it provides a secure connection between you and the remote host. Here's an example:

$ **sftp ftp.handsonhistory.com**

Connecting to ftp.handsonhistory.com

passwd: **********

sftp>

At this point you can begin an interactive FTP session. You can use `get` and `put` commands on files as you would using any FTP client, but with the comfort of knowing that you are working on a secure

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

connection.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip The sftp command, as with ssh and scp, requires that the SSH service be running on the server. If you can't connect to a FTP server using sftp, the SSH service may not be ava

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Using ssh, scp and sftp without passwords

For machines that you use a great deal, it is often helpful to set them up so that you do not have to use a password to log in. The following procedure shows you how to do that.
These steps will take you through setting up password-less authentication from one machine to another. In this example, the local user is named `chuckw` on a computer named `host1`. The remote user is also `chuckw` on a computer named `host2`.

1. Log in to the local computer (in this example, I log in as `chuckw` to `host1`).

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Only run step 2 once as local user on your local workstation. Do not run it again unless you lose your ssh keys. When configuring subsequent remote servers, skip right to step 3

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

2. Type the following to generate the ssh key:
3. $ **ssh-keygen -t dsa**
4. Accept the defaults by pressing Enter at each request.
5. Type the following to copy the key to the remote server (replace chuckw with the remote user name and host2 with the remote host name):
6. $ **cd ~/.ssh**
7. $ **scp id_dsa.pub chuckw@host2:/tmp**
8. chuckw@host2's password: ***********
9. Type the following to add the ssh key to the remote-user's authorization keys:
   $ ssh *chuck2@host2* 'cat /tmp/id_dsa.pub >> /home/*chuckw*/.ssh/authorized_keys2'

   Note Steps 3 and 4 will ask for passwords. This is okay.
10. Type the following to remove the key from the temporary directory:
11. $ **ssh chuckw@host2 /bin/rm /tmp/id_dsa.pub**

    Note Step 5 should not ask for a password.

It is important to note that once you have this working, it will work regardless of how many times the IP address changes on your local computer. IP address has nothing to do with this form of authentication.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Guarding Your Computer with PortSentry

While LogSentry gathers and sorts log messages that may represent attempts to break into your computer system, the PortSentry takes a more active approach to protecting your system from network intrusions. PortSentry can be installed and configured on a Red Hat Linux system to monitor selected TCP and UDP ports, and can then react to attempts to access these ports (presumably by people trying to break in) in ways that you choose.

PortSentry acts as a nice complement to LogSentry by actively looking for intrusion behavior on network ports. When PortSentry perceives an attack, it reacts to the attack (in ways that you choose) and produces log messages about the activity that can be forwarded to the system administrator by LogSentry.

PortSentry operates in several different modes. Each of these modes can be applied to monitoring of TCP and UDP ports. The PortSentry modes include:

Basic: This is the mode PortSentry uses by default. Selected UDP and TCP ports in this mode are bound by PortSentry, giving the monitored ports the appearance of offering a service to the network.

Stealth: In this mode, PortSentry listens to the ports at the socket level instead of binding the ports. This mode can detect a variety of scan techniques (strobe-style, SYN, FIN, NULL, XMAS and UDP scans), but because it is more sensitive than basic mode, it is likely to produce more false alarms.

Advanced Stealth: This mode offers the same detection method as the regular stealth mode, but instead of monitoring only the selected ports, it monitors all ports below a selected number (port number 1023, by default). You can then exclude monitoring of particular ports. This mode is even more sensitive than Stealth mode and is, therefore, more likely to cause false alarms than regular stealth mode.

Note When a port is "bound" by PortSentry or any other network service daemon process, all requests that come to that port from the network are handled by the binding process. For example, when the httpd daemon binds to port 80, requests for Web services from the network are processed by httpd.

Besides selecting the PortSentry mode and the ports that are monitored, you can also choose the response to your computer being scanned. By default, PortSentry can log intrusion attempts and block access to the intruder. PortSentry also offers ways of using other tools to respond to intrusions, including firewall rules, route changes, and host denial configuration.

Downloading and installing PortSentry

The `portsentry` package is not included in the Red Hat Linux distribution. You can download the package from any Red Hat Linux FTP mirror site or from the `rpmfind.net` site. The portsentry package used in this chapter is portsentry-1.0-11.i386.rpm (the last version published by Red Hat as part of the 7.1 PowerTools). After PortSentry is downloaded, run the following command from the directory you downloaded it to:

rpm -i portsentry*

The installed `portsentry` package consists of several configuration files (in the `/etc/portsentry` directory), the `portsentry` start-up script (`/etc/init.d/portsentry`), and the `portsentry` command (in `/usr/sbin`). There are also README files of interest in the `/usr/share/doc/portsentry*` directory.

Using PortSentry as is

As with LogSentry, you don't need to do anything to get PortSentry to work after it is installed. By default, here is what PortSentry does when you install the portsentry package:

The `/etc/init.d/portsentry` start-up script runs automatically when you boot to run levels 3, 4, or 5 (levels 3 and 5 are most commonly used).

The following port numbers are configured to be monitored by PortSentry in basic mode:

TCP: 1, 11, 15, 143, 540, 635, 1080, 1524, 2000, 5742, 6667, 12345, 12346, 20034, 31337, 32771, 32772, 32773, 32774, 40421, 49724, 54320

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

UDP: 1, 513, 635, 640, 641, 700, 32770, 32771, 32772, 32773, 32774, 31337, 54321
In response to attacks (represented by scans of the ports being monitored), all further attempts to connect to any services for the protocol (TCP or UDP) will be blocked.

The computers that are blocked from accessing your system are listed in either the `portsentry.blocked.tcp` or `portsentry.blocked.udp` files (in the `/var/portsentry` directory), depending on which protocol was scanned (TCP or UPD). Removing entries from these files to restore access to blocked computers.

Configuring PortSentry

Chances are that you will want to make some changes to the way that PortSentry runs. To change how PortSentry behaves, modify the `/etc/portsentry/portsentry.conf` file. In that file, you can choose which ports to monitor, the mode in which to monitor, and the responses to take when a scan is detected. The responses can include:

Blocking access by the remote computer
Rerouting messages from the remote computer to a dead host
Adding a firewall rule to drop packets from the remote computer

The other file you may want to change is the `/etc/portsentry/portsentry.modes`. The `portsentry.modes` file simply contains the modes that `PortSentry` can be run in.

Changing the portsentry.conf file

To edit the `portsentry.conf` file, as `root` user, open the `/etc/portsentry/portsentry.conf` file using any text editor. The following sections describe the information that can be changed in that file.

Selecting ports

The `portsentry.conf` file defines which ports are monitored in basic and stealth modes. By default, only basic TCP and UDP modes are active, so only those ports are monitored (unless you change to one of the stealth modes). The `TCP_PORTS` and `UDP_PORTS` options define which ports are monitored. Here is how they appear in the `portsentry.conf` file:

```
TCP_PORTS="1,11,15,143,540,635,1080,1524,2000,5742,6667,12345,12346,200
  34,31337,32771,32772,32773,32774,40421,49724,54320"
UDP_PORTS="1,513,635,640,641,700,32770,32771,32772,32773,32774,31337,
  54321"
```

Unless you are a TCP/IP expert, you're probably wondering what services these ports represent. The Internet Assigned Numbers Authority (IANA) assigns services to UDP and TCP ports. You can see these assignments at the following Web address:

www.iana.org/assignments/port-numbers

Network services in Red Hat Linux (as well as other Linux/UNIX systems) obtain port number assignments from the `/etc/services` file. So, in general, you can simply check the `/etc/services` file to find out most of the services that are assigned to ports being scanned.

Ports assigned for monitoring are chosen based on a couple of different criteria. Lower port numbers (1, 11, 15, etc.) are chosen to catch port scanners that begin at port 1 and scan through a few hundred ports. If the scanner is blocked after accessing port 1, it won't be able to get information about any other ports that may be open on your computer. Another criterion is to include ports that are checked specifically by intruders because those services may be vulnerable to attack. They include the `systat` (port 11) and `netstat` (port 15) services.

You will want to remove ports from the list in the `portsentry.conf` file if you are actually running the service assigned to that port. On the other hand, you may want to add ports to the list if you are paranoid about attacks and you want a bit more coverage. The `portsentry.conf` file contains some examples that you can uncomment (remove the # sign) so that more ports are monitored.

If you change from basic to stealth scans (as described in the "Changing the portsentry.modes file" section, later in this chapter), the ports that are monitored are those defined by the `ADVANCED_PORTS_TCP` and `ADVANCED_PORTS_UDP` options. Here is how those two options are set by default:

```
ADVANCED_PORTS_TCP="1023"
ADVANCED_PORTS_UDP="1023"
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

The two preceding entries indicate that all ports from 1 to 1023 are monitored. Monitoring higher port numbers can result in many more false alarms, so this practice is not recommended. If you find that PortSentry is being tripped accidentally, you may want to exclude the ports being tripped by using the `ADVANCED_EXCLUDE_TCP` and `ADVANCED_EXCLUDE_UDP` options. The following example shows how these two values are set by default:

ADVANCED_EXCLUDE_TCP="111,113,139"

ADVANCED_EXCLUDE_UDP="520,138,137,67"

By default, `ident` and NetBIOS services for TCP (ports 111, 113, and 139) and `route`, NetBIOS, and `Bootp` broadcasts for UDP (ports 520, 138, 127, and 67) are excluded from the advanced scan. (The exclusion is because a remote computer may hit these ports without representing any misuse.) If you are running in stealth mode, you should likewise exclude any services that you are running on your system by adding their port numbers to this list.

Identifying configuration files

Besides the portsentry.conf file, there are several other configuration files used by PortSentry. You can identify the locations of these other files within the portsentry.conf file. Here are how those files are defined:

Hosts to ignore

IGNORE_FILE="/etc/portsentry/portsentry.ignore"

Hosts that have been denied (running history)

HISTORY_FILE="/var/portsentry/portsentry.history"

Hosts denied this session only (temporary until next restart)

BLOCKED_FILE="/var/portsentry/portsentry.blocked"

Chances are that you will not want to move the location of these configuration files. Here are some descriptions of what these files are used for:

The `portsentry.ignore` file contains a list of all IP addresses that you do not want blocked (even if they improperly try to access ports on your computer). By default, all IP addresses assigned to the local computer are added to this file. You can add IP addresses of trusted computers, if you like.

The `portsentry.history` file contains a list of IP addresses for computers that have been blocked from accessing your computer.

The `portsentry.blocked.*` files contain a list of computers that have been blocked from accessing your computer during the current session. The `portsentry.blocked.tcp` file contains IP addresses of computers that have improperly scanned TCP ports on your computer. Addresses of computers that have been blocked after scanning UDP ports are contained in the `portsentry.blocked.udp` file.

Access to ports on your computer is only blocked during the current session (that is, until the next reboot or restart of PortSentry). So, to more permanently exclude remote computers, you should impose other restrictions (such as by using the `/etc/hosts.deny` file, a firewall command, or a reroute to a dead host). These methods are described later in this chapter.

Choosing responses

Someone scanning a port is like them checking a door in your house to see if it is locked. In most cases, it indicates that someone is checking your system for weaknesses. That is why, when another computer scans your ports, the default response from PortSentry is to block further access from the other computer to your computer for the duration of the current session. No action is taken to permanently block access from that computer. The `BLOCK_UDP` and `BLOCK_TCP` options in the `portsentry.conf` file set which type of automatic response is taken when ports are scanned. Here is how these options are set by default:

BLOCK_UDP="2"

BLOCK_TCP="2"

The value in quotation marks determines how PortSentry reacts to a scan of your ports by another

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

computer. The following list describes each of these values.

A value of "2" (the default value) causes access to be temporarily blocked to services for the scanned protocol (TCP or UDP) and for the action to be logged. Also, if any commands were defined to be run by a `KILL_RUN_CMD` option, that command is then run. (This option is not configured by default.)
A value of "0" causes port scans to be logged, but not blocked.
A value of "1" causes the `KILL_ROUTE` and `KILL_HOSTS_DENY` options to be run. (See the following list for descriptions of these options.) By default, further requests from the remote computer will be rerouted to a dead host, and the remote host's IP address will be added to the `/etc/hosts.deny` file, thereby denying access to network services.

Following are some suggestions on options you can use to can change the responses to your ports being scanned:

`KILL_ROUTE:` This option runs the `/sbin/route` command to reroute requests from the remote computer to a dead host. By default, this option is set to the following value, which effectively drops the request from the remote computer:

```
KILL_ROUTE="/sbin/route add -host $TARGET$ gw 127.0.0.1"
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Instead of rerouting IP packets from the remote host, you can use firewall rules to deny access. If you use ipchains firewalls, uncomment the following line to deny access from the

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

KILL_ROUTE="/sbin/ipchains -I input -s $TARGET$ -j DENY -l"
This `ipchains` rule would deny (in other words, drop) all packets from the remote computer. To make this action permanent, you could add the `ipchains` options (from the `-I` to the end of the line) to the `/etc/sysconfig/ipchains` file, replacing the `$TARGET$` with the actual IP address of the computer you want to deny access to.

`KILL_HOSTS_DENY`: This option is used to deny requests for any network services that are protected by TCP wrappers. This option is set by default as follows:

KILL_HOSTS_DENY="ALL: $TARGET$"
With the preceding option set, `$TARGET$` is replaced by the IP address of the intruding remote computer and the line in quotes is added to the `/etc/hosts.deny` file. For example, if the remote computer's IP address were 10.0.0.59, the line that appears in `/etc/hosts.deny` would be:

ALL: 10.0.0.59

`KILL_RUN_CMD`: Instead of using firewalls, rerouting, or TCP wrappers to deny an intruding computer from accessing your computer, you can choose any command you like in response. With the `BLOCK_TCP` and `BLOCK_UDP` options set to "2", the value of `KILL_RUN_CMD` is run in response to a scan of your monitored ports.

The value of `KILL_RUN_CMD` should be the full path to the script you want to run, plus any options. To include the IP address of the remote computer or the port number that was scanned, you could include the `$TARGET$` or `$PORT$` variables, respectively. Here is how the example appears that you would want to modify:

KILL_RUN_CMD="/some/path/here/script $TARGET$ $PORT$"

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution Do not use any KILL_RUN_CMD to retaliate against the intruding remote host. Firstly, it is quite possible that the computer that is scanning your ports has itself been cracked and

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

PORT_BANNER: You can send a message to the person who sets off the PortSentry monitor by setting the `PORT_BANNER` option. By default, no message is defined. However, you can uncomment the following line to use that message. (An abusive message is not recommended.)

PORT_BANNER="** UNAUTHORIZED ACCESS PROHIBITED ***

YOUR CONNECTION ATTEMPT HAS BEEN LOGGED. GO AWAY."

The number of scans from an intruding computer that PortSentry will accept before setting off the responses described above can be set by using the `SCAN_TRIGGER` option. By default, that option is set as follows:

SCAN_TRIGGER="0"

The "0" value means that you won't accept any scans from an intruding system. In other words, the first scan will trip the PortSentry monitor. You can increase this value to be tolerant of one or more errant scans (though you probably won't want to).

Changing the portsentry.modes file

The `/etc/portsentry/portsentry.modes` file defines the modes in which the PortSentry command is run at boot time. Here is how that file appears by default:

tcp

udp

#stcp

#sudp

#atcp

#audp

The `tcp` and `udp` options are the basic PortSentry modes for the TCP and UDP services, respectively. Your other choices of options include *stealth TCP* (`stcp`) and *advanced stealth TCP* (`atcp`) and *stealth UDP* (`sudp`) and *advanced stealth UDP* (`audp`). Only run one TCP service and one UDP service. So, if you uncomment a stealth or advanced stealth service, be sure to add a comment in front of the appropriate basic service.

To activate the new services, you would then execute the following command:

/etc/init.d/portsentry restart

The new PortSentry modes will take effect immediately. Those new modes will also be in effect when your computer reboots.

Testing PortSentry

You can test that your ports are properly protected in different ways. What you want to do is run a program that a potential intruder would run and see if it trips the appropriate response from PortSentry. For example, you could use a port scanner to see how your ports appear to the outside world. You could also use a command, such as `telnet`, to try and set off a particular port to see if PortSentry catches it.

`nmap` is a popular tool for scanning TCP and UDP ports. You can give the `nmap` command a host name or IP address, and it will scan about 1500 ports on computer to see which ports are open (and presumably offering services that could potentially be cracked).

An RPM of `nmap` is also available. You can download the `nmap-frontend` package, which contains a simple graphical interface to `nmap` called `xnmap`. I suggest that you install the packages on the system running PortSentry as well as on another system on your LAN (if one is available). Then run the following procedure on the PortSentry system to test it:

1. If PortSentry is running, shut it down by typing the following:
2. # **/etc/init.d/portsentry stop**
3. Type the following `nmap` commands to see which ports are open on the local system:
4. # **nmap -sS -O 127.0.0.1**
5. # **nmap -sU -O 127.0.0.1**

The output shows you which ports are currently offering services on your computer for TCP

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

and UDP protocols, respectively.

6. If there are any services that you don't want open, you should turn off those services by using `chkconfig service off` (replacing *service* with the service name), by editing the configuration file in the `/etc/xinetd.d` directory that represents the service and changing `disable = no` to `disable = yes`, or by changing your firewall setup.
7. If there are services that you want to be available from your computer, make sure that the port numbers representing those services are not being monitored by PortSentry. Remove the port number from the `TCP_PORTS` and/or `UDP_PORTS` options in the `/etc/portsentry/portsentry.conf` file, or PortSentry will report that there is a possible stealth scan on the port.
8. Restart PortSentry as follows:
9. # **/etc/init.d/portsentry start**
10. Run `nmap` again, as described previously. The ports offering legitimate services, as well as the ports being monitored by PortSentry, should all appear to be open.
11. Check the `/var/log/messages` file to make sure that PortSentry is not trying to monitor any ports on which you are offering services.

When you have determined that PortSentry is set up the way you would like it to be, run the `nmap` command from another computer on your network. This time, replace 127.0.0.1 with the name or IP address of the PortSentry computer. If everything is working, the first port that the remote computer scans on your PortSentry computer should block all subsequent scans.

Tracking PortSentry intrusions

Besides taking action against intruders, PortSentry logs its activities using the `syslog` utility. PortSentry's start-up, shutdown, and scan-detection activities are logged to your `/var/log/messages` file. The following are examples of PortSentry output in that file.

portsentry[13259]: adminalert: Psionic PortSentry 1.0 is starting.

portsentry[13260]: adminalert: Going into listen mode on TCP port: 1

portsentry[13260]: adminalert: Going into listen mode on TCP port: 11

.

.

.

portsentry[13260]: adminalert: PortSentry is active and listening.

portsentry[]: attackalert:Connect from host:10.0.0.4 to TCP port: 31337

portsentry[]: attackalert: Connect from host: 10.0.0.4 to TCP port: 11

portsentry[]: attackalert: Host: 10.0.0.4 is already blocked. Ignoring

.

.

.

portsentry[13371]: securityalert: Psionic PortSentry is shutting down

portsentry[13371]: adminalert: Psionic PortSentry is shutting down

The first part of the output shows PortSentry starting up. As PortSentry begins listening to each port, that port is noted in a separate log message. The next messages show the computer being scanned. Someone from host 10.0.0.4 ran `nmap` to scan the ports on the computer. PortSentry caught the scan of port 31337 and blocked attempts to scan other ports.

Finally, the last set of messages shows PortSentry being shut down. This is a security alert because someone besides you could shut down PortSentry to hide that they had broken in.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you have been running the LogSentry package (described earlier), these messages show up in the e-mail messages you receive each hour from Logcheck.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Restoring access

If access was cut off to a computer that you wanted to have access, there are several things you can check to correct that problem:

`/etc/hosts.deny:` See if the computer's IP address was mistakenly added to this file. This would cause network services to be denied to the host at that IP address.

`/var/portsentry/portsentry.blocked:` Check that an entry for the computer's IP address wasn't added to the `portsentry.blocked.udp` or `portsentry.blocked.tcp` files.

`route:` Run the `/sbin/route` command to see if messages from the computer are being rerouted to a dead host (probably the localhost).

`ipchains:` Run the `ipchains -L` command to see if a new firewall was created to block access from the computer. (If you had created an iptables rule instead, type **`iptables -L`** to see the new firewall entry.)

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Tip To make sure that access isn't cut off again, you can add the IP address of the remote computer to the /etc/portsentry/portsentry.ignore file. Future improper scans or requests for ser

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Part IV: Red Hat Linux Network and Server Setup

Chapter List

Chapter 15: Setting Up a Local Area Network
Chapter 16: Connecting to the Internet
Chapter 17: Setting Up a Print Server
Chapter 18: Setting Up a File Server
Chapter 19: Setting Up a Mail Server
Chapter 20: Setting Up an FTP Server
Chapter 21: Setting Up a Web Server
Chapter 22: Setting Up a News Server
Chapter 23: Setting Up Boot Servers: DHCP and NIS
Chapter 24: Setting Up a MySQL Database Server
Chapter 25: Making Servers Public with DNS
Chapter 26: Using Linux Servers from a Mac
Appendix A: What's on the CD-ROMs
Appendix B: Red Hat Linux RPMs
Appendix C: Running Network Services

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Understanding Local Area Networks

Connecting the computers in your organization via a LAN can save you a lot of time and money. The amount of money you put into networking hardware, even in a small configuration (less than five or six users), can save you from buying multiple printers, backup media, and other hardware. Add a single, shared Internet connection and you no longer need multiple modems and Internet accounts.

With a LAN, you don't have to run down the hall anymore with your file on a disk to print it on your friend's printer. Information that had to wait for the mailroom to make the rounds can be sent in an instant to anyone (or everyone) on your LAN.

With a LAN, you begin to open the greatest potential of Red Hat Linux — its ability to act as a server on a network. Because Red Hat Linux is more robust and feature-rich than other computing systems (certainly for the price), adding it to your LAN can provide a focal point to workstations that could use Red Hat Linux as a file server, a mail server, a print server, or a news server. (Those features are described later in this book.)

Creating and configuring a LAN consists of these steps:

1. **Setting up LAN hardware** — This entails choosing a network topology, purchasing the equipment you need, and installing it (adding cards and connecting wires or using wireless antennas).
2. **Configuring TCP/IP** — To use most of the networking applications and tools that come with Red Hat Linux, you must have TCP/IP configured. TCP/IP lets you communicate not only with computers on your LAN, but to any computers that you can reach on your LAN, modem, or other network connection (particularly to the Internet).

Setting up LAN hardware

Even with a simple LAN, you must make some decisions about network topology (that is, how computers are connected). You must also make some decisions about network equipment (network interface cards, wires, hubs, and so on).

LAN topologies

Most small office and home LANs connect computers together in one of the following topologies:

1. **Star topology —** The star topology is by far the most popular LAN topology. In this arrangement, each computer contains a Network Interface Card (NIC) that connects with a cable to a central hub. The cabling is typically Category 5 (twisted pair) wiring with RJ-45 connectors. Other equipment, such as printers and fax machines, can also be connected to the hub in a star topology. Figure 15-1 is an example of a star topology.

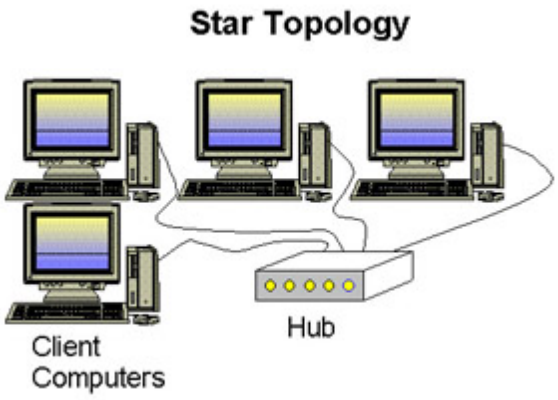

Figure 15-1: In a star topology, machines on the network connect to a central hub.

2. **Bus topology —** Instead of using hubs, the bus topology connects computers in a chain from one to the next. The cabling usually used is referred to as coaxial, or Thin Ethernet cable. A "T" connector attaches to each computer's NIC, then to two adjacent computers in the chain. At the two ends of the chain, the T connectors

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

are terminated. Figure 15-2 is an example of a bus topology.

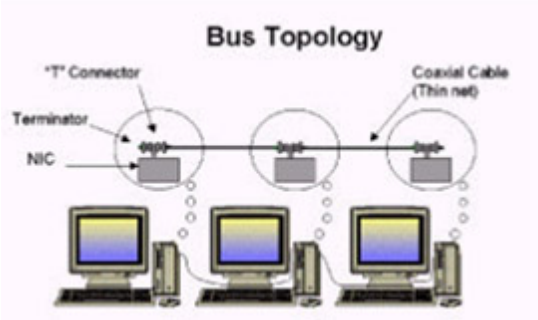

Figure 15-2: A bus topology chains computers together without using a hub.

3. **Ring topology —** This is a less popular topology than star and bus topologies. In a ring topology, computers connect to a ring of wires on which tokens are taken and passed by computers that want to send information on the network. This type of topology typically uses IBM's token ring protocols.

You can configure a wireless Ethernet LAN in several different topologies, depending on how you want to use the LAN. With a wireless LAN, each computer broadcasts in the air rather than across wires. Here are some examples of wireless topologies:

Wireless peer-to-peer — In this topology, frames of data are broadcast to all nodes within range, but are consumed only by the computers for which they are intended. This arrangement is useful if you are sharing file and print services among a group of client computers. Figure 15-3 shows an example of a peer-to-peer wireless LAN.

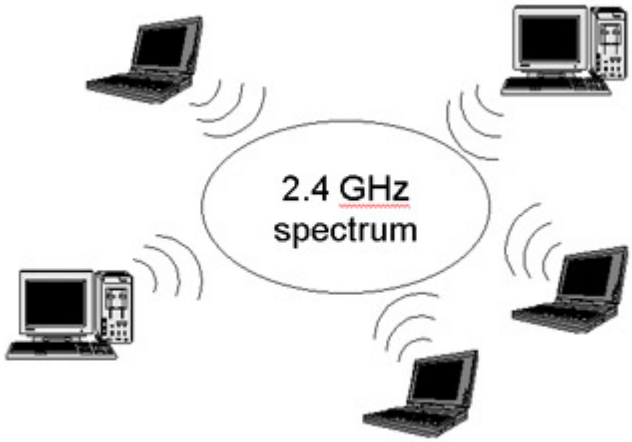

Figure 15-3: Wireless LANs can communicate as peers by broadcasting data.

Wireless access point — A wireless interface can act as an access point for one or more wireless clients. Clients can be configured to communicate directly with the access point, instead of with every client that is within range. This arrangement is useful for point-to-point connections between two buildings, where the access point is acting as a gateway to the Internet or, for example, a campus intranet. Figure 15-4 depicts a point-to-point wireless LAN.

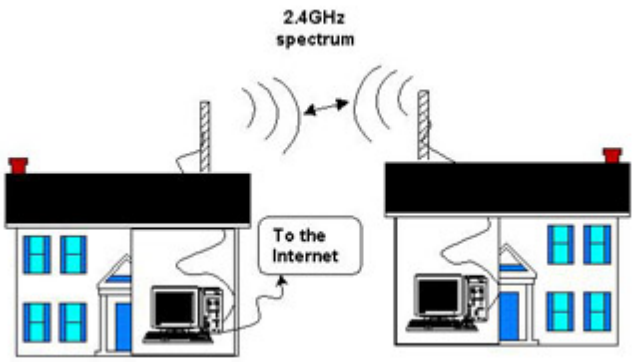

Figure 15-4: Wireless communication can go through an access point.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Multiple wireless access points and roaming — Each wireless network can be configured as a separate cell. Multiple cells can be joined together in what is called a *Managed Wireless LAN*. In this arrangement, each cell's access point acts as a bridge (in fact, its actions are referred to as *bridging*) by passing all data from the cell to other cells without changing any node's MAC address. In other words, the managed wireless LAN makes the fact that there are multiple wireless cells invisible to the clients on those cells. This arrangement allows users to roam among cells as they come in and out of range without losing continuity of communication.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Although Linux does not support any wireless LAN cards that can act as a bridging wireless access point, you can have the Linux system act as a client to a bridging access point

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

For our purposes, we focus on star (wired) and peer-to-peer (wireless) topologies. Common to both of these topologies is the protocol used to send data over those wired and wireless media — the Ethernet protocol.

LAN equipment

The equipment that you need to connect your LAN can include some or all of the following:

Cross-Reference For a complete description of wireless hardware, see the "Choosing wireless hardware" section later in this chapter.

Network Interface Card (NIC) — Typically, one of these cards goes into a slot in each computer. For wired Ethernet networks, the cards can transmit data at 10 Mbps or 100 Mbps. Gigabit (1000 Mbps) NICs are also now available, but are quite a bit more expensive. An 802.11B wireless NIC card can operate at speeds of up to 11 Mbps, but are more expensive than a wired NIC card.

Cables — For star topologies, cables are referred to as *twisted-pair*. Category 5e wiring, which contains four twisted-pair sets per wire, is the most common type of wiring used for LANs today. A connector at each end of the cable is an RJ-45 plug, similar to those used on telephone cables. Ethernet interfaces are either 10Base-T (10 Mbps speeds) or 100Base-TX (100 Mbps speeds). These cables plug into the computer's NIC at one end and the hub at the other.

Figure 15-5 shows an example of a twisted-pair cable with an RJ-45 connector used for star topologies.

Figure 15-5: A star topology's twisted-pair cables have RJ-45 connectors (similar to telephone-cable connectors).

Hubs — With the star topology, a hub is typically used to connect the computers. Sometimes hubs are also referred to as repeaters or concentrators because they receive signals from the nodes connected to them and send the signals on to other nodes.

The questions you need to answer when choosing a hub are how many ports you need and how fast you want your network to go. If you need to expand in the future, hubs can be connected together. Most low-end hubs you can purchase today handle both 10 Mbps and 100 Mbps speeds.

Switches — A switch can be used instead of a hub. It lets you divide a LAN that is getting too large into segments that are more manageable. A switch can reduce network traffic by directing messages intended for a specific computer directly to that computer. This is as opposed to a hub, which broadcasts all data to all nodes. Because switches have come down so much in price, in most cases you should

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

just pay a few extra dollars and get a switch instead of a hub.

One piece of equipment that I won't go into yet is a router. A router is used to direct information from the LAN to other LANs or the Internet.

Cross-Reference Machines that carry out routing functions are described in Chapter 16.

LAN equipment setup

With an Ethernet NIC, appropriate cables, and a hub (or switch), you are ready to set up your wired Ethernet LAN. If you don't yet have an Ethernet card for your computer, refer to the sidebar "Choosing an Ethernet Card" for information on choosing a card that will work in Linux. The steps for setting up an Ethernet LAN are:

1. Power down each computer and physically install the NIC card (following the manufacturer's instructions).
2. Using cables appropriate for your NIC cards and hub, connect each NIC to the hub.
3. Power up each computer.
4. If Red Hat Linux is not installed yet, install the software and reboot (as instructed). Chapter 2 tells you how to configure your Ethernet card while installing Red Hat Linux.
5. If Red Hat Linux is already installed, refer to the "Configuring TCP/IP for your LAN" section for information on configuring your Ethernet cards.
6. When the system comes up, your Ethernet card and interface (`eth0`) should be ready to use. See the section "Checking your Ethernet connection" later in this chapter to learn how to determine whether your Ethernet connection is working.

Choosing an Ethernet Card

There are several ways to find out about supported Ethernet cards for Linux. Here are some suggestions:

The `/usr/src/linux*/Documentation/networking` directory contains files that describe Ethernet cards that are supported in the Linux kernel. You need the kernel-source package installed to see these files.

For laptops, PCMCIA cards supported in Red Hat Linux are defined in the `/etc/pcmcia/config` file. More than 100 PCMCIA Ethernet cards are listed, a handful of wireless PCMCIA LAN drivers, and a couple of Token Ring PCMCIA cards.

The Linux Ethernet-HOWTO (`www.tldp.org/HOWTO/Ethernet-HOWTO.html`) describes many older Ethernet cards that are supported.

Configuring TCP/IP for your LAN

When you install Red Hat Linux, you are given the opportunity to add your TCP/IP host name and IP address, as well as some other information, to your computer or choose to have that information automatically provided using DHCP. You also can set up a way to reach other computers on your LAN. That's typically done by adding computer names and IP addresses to your `/etc/hosts` file (as described here) or by using a DNS server.

Cross-Reference DNS is discussed in Chapter 16. Configuring your own DNS server is described in Chapter 25.

If you did not configure your LAN connection during installation of Red Hat Linux, you can do so at any time using the Network Configuration window (`neat` command). The IP address and host names can be assigned statically to an Ethernet interface or gotten dynamically at boot time from a DHCP server.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note A computer can have more than one IP address because it can have multiple network interfaces. Each network interface must have an IP address (even if the address is as

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To define your IP address for your Ethernet interface, follow this procedure:

1. Start the Network Configuration. From the Red Hat menu, click System Settings → Network or, as root user from a Terminal window, type `neat`. (If prompted, type the root password.) The Network Configuration window appears.
2. Click the Devices tab. A listing of your existing network interfaces appears.
3. Double-click the eth0 interface (representing your first Ethernet card). A pop-up window appears, enabling you to configure your eth0 interface. Figure 15-6 shows the Network Configuration window and the pop-up Ethernet Device window configuring eth0.

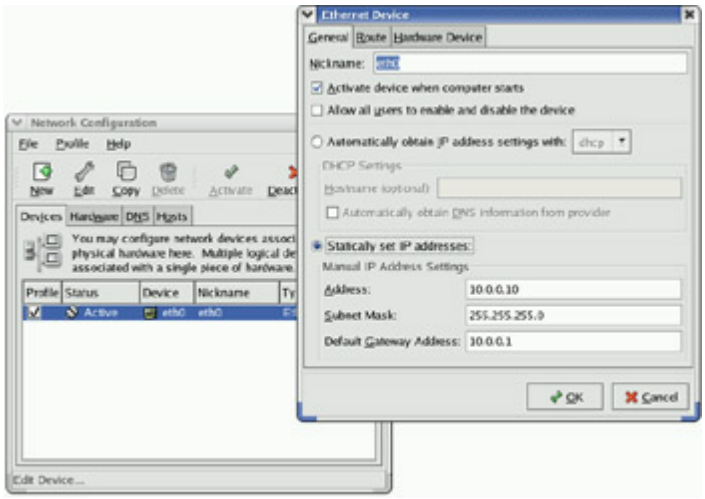

Figure 15-6: Configure your LAN interface using the Network Configuration window.

4. On the Ethernet Devices window that appears, you can enter the following information:

   Activate device when computer starts: Check here to have eth0 start at boot time.

   Allow all users to enable and disable the device: Check to let non-root users enable and disable the network interface.

5. On the same window, you must choose whether to get your IP addresses from another computer at boot time or enter the adresses yourself:

   Automatically obtain IP address settings with: Select this box if you have a DHCP or BOOTP server on the network from which you can obtain your computer's IP address, netmask, and gateway. DHCP is recommended if you have more than just a couple of computers on your LAN. (See Chapter 23 for how to set up a DHCP server.) You can, optionally, set your own host name, which can be just a name (such as `jukebox`) or a fully qualified domain name (such as `jukebox.linuxtoys.net`).

   Statically set IP addresses: If there is no DHCP, or other boot server, on your LAN, you can add necessary IP address information statically by selecting this option and adding the following information:

   Address: Type the IP address of this computer into the Address box. This number must be unique on your network. For your private LAN, you can use private IP addresses (see the section, "Understanding IP addresses" later in this chapter).

   Subnet Mask: Enter the netmask to indicate what part of the IP address represents the network. (Netmask is described later in this chapter.)

   Default Gateway Address: If a computer or router connected to your LAN is providing routing functions to the Internet or other network, type the IP address of the computer into this box. (Chapter 16 describes how to use NAT or IP masquerading and use Red Hat Linux as a router.)

6. Click OK in the Ethernet Device window to save the configuration and close the window.
7. Click File → Save to save the information you entered.
8. Click Activate in the Network Configuration window to start your connection to the LAN.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Identifying other computers (hosts and DNS)

Each time you use a name to identify a computer, as when browsing the Web or using an e-mail address, the computer name must be translated into an IP address. To resolve names to IP addresses, Red Hat Linux goes through a search order (based on the contents of three files in `/etc`: `resolv.conf`, `nsswitch.conf`, and `host.conf`). By default, it checks:

Host names you add yourself (which end up in the `/etc/hosts` file).
Hosts available via NIS (if an NIS server is configured as described in Chapter 23).
Hosts names available via DNS.

You can use the Network Configuration window to add:

Host names and IP addresses. You might do this to identify hosts on your LAN that are not configured on a DNS server.
DNS search path. By adding domain names to a search path (such as `linuxtoys.net`), you can browse to a site by its host name (such as `jukebox`), and have Linux search the domains you added to the search path to find the hostyou are looking for (such as`jukebox.linuxtoys.net`).
DNS name servers. A DNS server can resolve addresses for the domains it serves and contact other DNS servers to get addresses for all other DNS domains.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you are configuring a DNS server, you can use that server to centrally store names and IP addresses for your LAN. This saves you the trouble of updating every computer's /et

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

To add host names, IP addresses, search paths, and DNS servers, do the following:

Start the Network Configuration. As root user from a Terminal window, type `neat&` or from the Red Hat menu, click System Settings → Network. The Network Configuration window appears.
Click the Hosts tab. A list of IP addresses, host names, and aliases appears.
Click New. A pop-up window appears asking you to add the IP address, host name, and aliases for a host that you can reach on your network. Figure 15-7 shows the Network Configuration window and the pop-up window for adding a host.

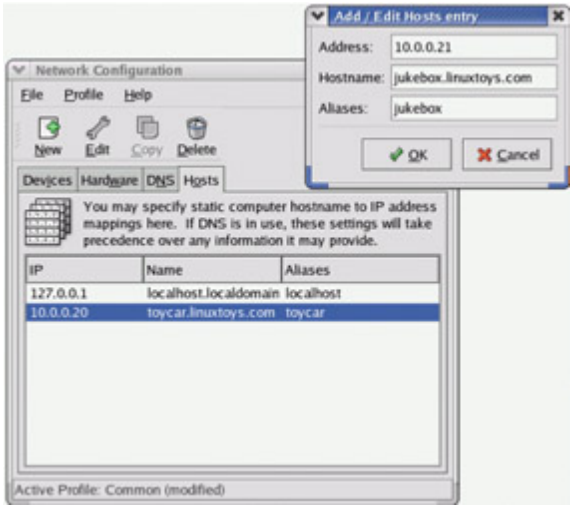

Figure 15-7: Add hosts to /etc/hosts using the Network Configuration window.

Type in the IP address number, host name, and, optionally, the host alias.
Click OK.
Repeat this process until you have added every computer on your LAN.
Click the DNS tab.
Type the IP address of the computers that serve as your Primary and Secondary DNS servers. You get these IP addresses from your ISP or, if you created your own DNS server, you can enter that server's IP address.
Type the name of the domain (probably the name of your local domain) to be searched for host names into the DNS Search Path box.
Click File → Save to save the changes.
Click File → Quit to exit.

Now, when you use programs such as `ftp`, `ssh`, or other TCP/IP utilities, you can use any host name that is identified on your local computer, exists in your search path domain, or can be resolved from the public Internet DNS servers. (Strictly speaking, you don't have to set up your `/etc/hosts` file. You could use IP addresses as arguments to TCP/IP commands. But names are easier to work with.)

Adding Windows computers to your LAN

It is likely that you have other types of computers on your LAN in addition to those running Red Hat Linux systems (at least for a few more years). The following are general steps for adding your Windows computers to the Ethernet LAN we just created:

1. Power down your computer and install an Ethernet card. (Most PC Ethernet cards will run on Windows.)
2. Connect an Ethernet cable from the card to your hub.
3. Reboot your computer. If your card is detected, Windows will either automatically install a driver or ask you to insert a disk that comes with the card to install the driver.
4. Open the window to configure networking. (Start → Settings → Control Panel; then double-click the Network icon. If you have Windows XP, you also need to click "Set up or change your Internet connection."). A window to change network properties appears.
5. What you do next depends on the version of Windows you are running:
   For Windows 98:
   1. Find the Ethernet card you have just installed in the list and select it.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

2. Click Add. The Select Network Component Type pop-up window appears.
3. Double-click Protocol. The Select Network Protocol window appears.
4. Click Microsoft, and then double-click TCP/IP. A new entry should appear in your Network window that looks similar to the following, depending on your card:

   TCP/IP -> 3Com Etherlink III ISA
5. Double-click on that new entry. The TCP/IP Properties window should appear, similar to the one in Figure 15-8 for Windows XP.

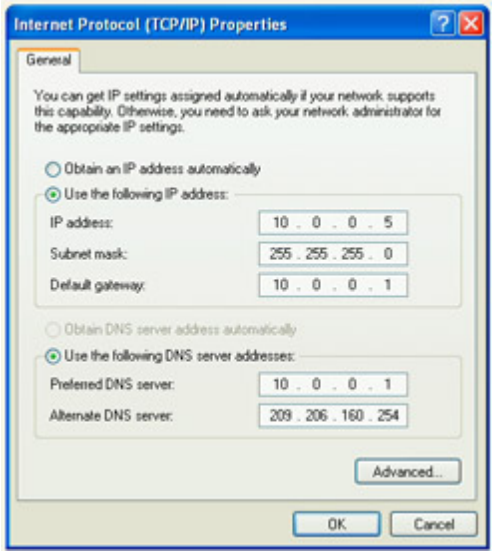

Figure 15-8: Configure TCP/IP on Windows XP for your Ethernet LAN.

For Windows 2000 or XP:

6. Click "Switch to classic view."
7. Double-click the Network connections.
8. Double-click Local Area Connection. Local Area Connection Status window appears.
9. Click Properties. The Local Area Connection Properties window appears.
10. Select Internet Protocol (TCP/IP), and click the Properties button. The Internet Protocol (TCP/IP) Properties window appears as shown in Figure 15-8.

6. Click either "Use the following IP address" or "Obtain an IP address."

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note If you are using a DHCP server to assign IP addresses, click "Obtain an IP address automatically" instead. See Chapter 23 for information on onsetting up Red Hat Linux as a DHCP server.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

7. Add the IP address, Subnet mask, and Default Gateway for this computer.
8. Add the IP addresses of up to two DNS servers.
9. Click OK. You may need to reboot Windows for the settings to take effect.

At this point, your Windows computer knows to listen on the network (via its Ethernet card) for messages addressed to the IP address you have just entered.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Setting Up a Wireless LAN

Sometimes it's not convenient to run wires to all the computers on your network. Pulling Ethernet cables through existing walls can be a pain. Dragging wires into your garden so you can sit in a lounge chair and surf the Internet can ruin the ambiance. In many cases, a wireless LAN is an economical solution.

Although you can use wireless LAN cards with other computer systems, you may want to use Linux systems for one or more nodes in your wireless network. For example, the features in Linux can eliminate the need to buy other types of equipment. Some additional features that make Linux a valuable asset on a wireless LAN include:

Internet access — You don't need a separate router or gateway machine to attach your wireless LAN to the Internet. Having wired and wireless Ethernet LAN cards on a Linux system enables your wireless clients to access the Internet through your Linux system.

Firewall — Owing to some inherent security weaknesses with wireless encryption protocols, you may want to add an extra measure of security to your network by configuring firewalls. With a Red Hat Linux firewall (ipchains or iptables) at the boundaries between your wireless LAN and your larger network, you still have a measure of protection for your larger network if someone cracks your wireless LAN.

Monitoring and logging — All the tools you use for monitoring and logging activity on your wired networks in Red Hat Linux are also available for your wireless network.

This chapter describes how to use wireless LAN equipment on computers running Red Hat Linux to create a wireless Ethernet LAN. It focuses on configuring two Red Hat Linux systems for wireless communication; however, once you configure these nodes, you can add Windows, Linux, or other types of systems to your wireless LAN by installing compatible wireless cards on each system.

Understanding wireless networks

Wireless LANs are most appropriate in environments where wires are impractical. Despite some challenges such as security and interference, a wireless LAN provides these advantages:

You don't have to run wires in places that are hard to reach. In many cases, a single wireless LAN can extend your network throughout a building or to another building without the need for wires between each node.

For the price of a wireless card, you can save the expense of wires, hubs (the air is your hub), and wall repairs (to fix the holes from pulling wires through).

You can freely move computers around within the transmission range that your environment allows (distances being limited by such variables as antenna power, obstacles, and rates of transmission).

Although several different wireless networking standards exist, this chapter focuses on the installation of relatively low-cost, standard IEEE 802.11b wireless-networking equipment. An 802.11b wireless network uses space in the spectrum available to the public (in other words, you use space in the air for which no special license is required). The 802.11b standard is often referred to as the *Wi-Fi,* or *Wireless Fidelity*, standard.

Note If you are interested in exploring 802.11a or 802.11g wireless LAN cards, I recommend you check out the MADWIFI project (http://sourceforge.net/projects/madwifi). The project is developing Linux drivers for wireless LAN cards (802.11a/b/g) that contain Atheros chip sets.

An 802.11b network is characterized by the following:

It provides transmission rates of up to 11 Mbps. Transmission rates can also be set (or auto-detected) to 5.5, 2, and 1 Mbps.

It uses the 2.4 GHz band of the spectrum. Microwave ovens and some high-end mobile phones also use this band. (Check local regulations if you are setting up an 802.11b network outside the United States.) To reduce congestion, 14 separate channels have been made available within the 2.4 GHz range.

It allows transmission over distances as short as a desktop away to as long as

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

several miles away (using special antennas). Greater distances can be gained at lower transmission speeds.

Makes connections between multiple clients or clients and a base station (usually referred to as an *access point*). On the clients, the wireless LAN cards run in Ad-hoc mode, while the base station uses Managed mode.

Note The Orinoco card, as well as other wireless-network cards supported in Linux, cannot act as an access point because it does not do bridging. Bridging allows a node to receive a frame from one node and forward it to another node without changing the first node's MAC address. (The MAC address uniquely identifies a network card.) A wireless LAN card in Linux, however, can communicate with an access point by running in Managed mode and indicating the MAC address of the access point.

Other 802.11 standards exist (such as 802.11a, which can operate at higher speeds), but for the most part wireless-equipment manufacturers have rallied around the 802.11b standard. Wireless cards and other equipment certified Wi-Fi (802.11b)–compatible by the Wireless Ethernet Compatibility Alliance (WECA) should be able to communicate with each other.

Cross-Reference To see a complete list of Wi-Fi–certified products, visit the WECA Web site (www.wirelessethernet.org, then click Certified Products). Although these products should be able to communicate with each other, they do not all have drivers that lets them work in Linux.

After your wireless network has been configured, you can use the wireless connections as you would a regular wired Ethernet connection. For example, you can configure TCP/IP on top of your wireless network so that it acts as a gateway to your network's Internet connection. If you are using Linux as a wireless network client as well, you can take full advantage of firewall, masquerading, network proxy, or other networking features to protect and make full use of your wireless network.

Choosing wireless hardware

To get started with a wireless Linux LAN, you need at least two computers and two wireless LAN cards. The wireless LAN cards described in this chapter are PCMCIA-type cards that you insert into those credit-card–sized slots on laptop computers. If ISA or PCI slots alone are available on your computers, you will need to add an adapter card. You may also want (or need) to add indoor or outdoor antennas to your wireless network.

Selecting wireless cards

Not all wireless LAN cards that you can purchase today will work with Linux. When you select a card, make sure that a Linux driver has been created for the card. You will also want to look for several other features:

Card type — Most wireless LAN cards are in PCMCIA (PC-card) form. To use these cards in desktop computers you will probably need an ISA or PCI adapter card.

Cost — If you shop around, you can find wireless LAN cards that cost between $29 and $100 (and the prices will probably have gone down by the time you read this). More expensive cards may include external antennas or better encryption (128-bit as opposed to 64-bit).

External connector — Different wireless LAN cards have different types of connectors, which can lock you into buying antennas from the manufacturer of the LAN card if you are communicating outside of a small area. People who know such things tell me that it is possible to take apart the cards and hack together your own antenna. However, because I don't want to be responsible for wrecking your card, and because there are legal issues related to antenna usage, I don't recommend this.

Configurability — If you really want to fine-tune your wireless LAN, find out how much control you have over configuring a card before you choose it. Descriptions of the `iwconfig` command later in this chapter will help you understand which wireless extensions you can manipulate. Refer to man pages for individual wireless-card drivers (for example, type `man wvlan_cs`), for information on specific parameters that you can change.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

I chose to use Orinoco wireless PC cards from Proxim Corporation to illustrate how to set up a wireless Linux LAN. Although other wireless cards are supported in Linux, Orinoco cards seem to have particularly good Linux drivers, and many people have reported success using them. The cards are also relatively inexpensive (about $60 to $90 at the time of this writing). Before Proxim acquired the Orinoco line from Agere Systems, Orinoco cards were referred to as *WaveLAN cards* and were produced by Lucent Technologies (as well as AT&T and NCR).

Other wireless cards supported in Red Hat Linux are listed in the PCMCIA Card Configuration Database file (`/etc/pcmcia/config`). Table 15-1 shows the wireless-network adapters listed in this file, along with the module required by each.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Just because you don't see your wireless card listed in Table 15-1 doesn't mean that it won't work. The same technology is often referred to by different names. Rather than try to

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 15-1: Supported Wireless Network Adapters and Modules

Wireless Network Adapter	Module
350 Series Wireless LAN Adapter (Cisco Systems)	airo_cs
Aironet PC4500 (Cisco Systems)	airo_cs
Aironet PC4800 (Cisco Systems)	airo_cs
AT&T WaveLAN Adapter	wavelan_cs
Cabletron RoamAbout 802.11 DS	wvlan_cs
Compaq WL100 11 Mbps Wireless Adapter	orinoco_cs
Digital RoamAbout/DS	wavelan_cs
ELSA AirLancer MC-11	wvlan_cs
Lucent Technologies WaveLAN Adapter	wavelan_cs
Intersil PRISM2 11 Mbps Wireless Adapter	wvlan_cs
Lucent Technologies WaveLAN/IEEE Adapter	orinoco_cs
MELCO WLI-PCM-L11	orinoco_cs
MELCO WLI-PCM-L11G	orinoco_cs
NCR WaveLAN Adapter	wavelan_cs
NCR WaveLAN/IEEE Adapter	wvlan_cs
Orinoco PC Cards	orinoco_cs
PLANEX GeoWave/GW-CF110	orinoco_cs
RayLink PC Card WLAN Adapter	ray_cs
Xircom CreditCard Netwave	netwave_cs
ZCOMAX AirRunner/XI-300	orinoco_cs

The Orinoco PC cards I purchased use the orinoco_cs module. The cards come in two types: Gold Label, which offers 128-bit WEP RC4 encryption, and Silver Label, which offers only 64-bit WEP RC4 encryption. The two cards differ in cost by about $10. Both cards:

- offer compatibility with earlier WaveLAN/IEEE products.
- can communicate with other 802.11b wireless LAN equipment that has been Wi-Fi–certified by WECA.
- enable you to select transmission rates (11, 5.5, 2, or 1 Mbps), select channels (in the 2.4 GHz range), and use power-management features.

Because antennas are built into Orinoco PC cards, you may not need an additional antenna for indoor office environments. As with any wireless card, the distances you can achieve between your wireless nodes depend on transmission rate (lower speeds go farther), receiver sensitivity, and the amount and type of obstacles. At the maximum transmission rate (11 Mbps), the Orinoco card gives you an estimated range of between 80 feet (25 meters) in a closed area and 525 feet (160 meters) in an open area with no additional antenna.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Selecting adapter cards

If the computers on your wireless network are all laptops with PCMCIA slots, you only need to plug in your wireless cards to get started. However, if you are using a desktop computer with only ISA and PCI slots, you will also need an adapter card.

Before you purchase a wireless LAN card for your desktop computer, make sure that you can get a compatible adapter card. Proxim Corp. offers both ISA and PCI Orinoco adapter cards. If you have an ISA slot available, you should get an ISA adapter. The ISA card supports the following I/O addresses: 3E2-3E1 (default) and 3E2-3E3.

The PCI adapter works on computers that have the following features:

- PCI 2.2 (or higher) BIOS support
- PC99 compliance
- PCI slots only

The PCI adapter will not work on some older computers that don't meet the BIOS specifications. To use the PCI adapter in Linux, you may have to perform additional configuration in the operating system.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note Although using the Orinoco IS A and PCI adapters with a standard Type II PC is not specifically supported, many Type II PC cards (besides your wireless cards) will work in those adapt

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Selecting antennas

If you are setting up your wireless LAN among several computers in close proximity to each other, you may not need an additional antenna. To deal with obstructions and longer distances, however, you can add indoor or outdoor antennas to your wireless hardware.

Again, because I have been discussing Orinoco wireless PC cards, I will illustrate different types of indoor and outdoor antennas that are compatible with those cards.

Using indoor antennas

The antennas that are built into wireless LAN cards often work well enough to enable communication among computers in an open area. Additional indoor antennas are useful if the direct line of sight between the wireless LAN cards is blocked. A computer may be locked in a storage closet or stuck under a desk. A pile of papers might inhibit transmission, or a sheet of metal might stop it dead. A small antenna that draws the transmission away from the card might be the answer to these problems.

While most wireless LAN cards don't require a completely unobstructed line of sight, an obstacle can certainly slow reception. To get around this problem, an antenna such as the Orinoco IEEE range-extender can plug directly into an Orinoco Gold or Silver wireless LAN card. A 1.5-meter extension cable can bring the signal out from behind a closed door or out on top of a desk. When you set up the antenna, it is recommended that it be:

placed in a central location.
mounted vertically.
located away from obstructions (metal surfaces in particular, and, to a lesser extent, solid objects such as concrete walls or stacks of papers).

Refer to the instructions that come with your antenna for specific guidelines regarding placing and mounting the antenna.

Using outdoor antennas

Choosing and setting up outdoor antennas for your wireless LAN can be more difficult and expensive than setting them up indoors. Once the outdoor antennas are in place, however, you can save money because you won't need multiple Internet access accounts (monthly fees, DSL/cable modems, and so on).

Although a complete description of the use of outside antennas with your wireless LAN is outside the scope of this chapter, here are some tips that will help you choose the best antennas for your wireless LAN.

Point-to-point versus multi-point — If you are creating a point-to-point link between two outdoor locations (for example, to share an Internet connection between two buildings), a directional antenna can help you achieve greater distance and transmission speeds. However, if your antenna is providing multi-point access for several other outdoor antennas or wireless clients (such as students working from laptops on the campus lawn), an omnidirectional antenna may be more appropriate.

Clearance — The clearer the line of sight between each outdoor antenna, the greater the distance and transmission speed you can achieve. Placing antennas at the highest possible points can prevent diminished performance caused by trees, cars, buildings, and other objects. The amount of distance between obstacles and the coverage area of your wireless transmission is referred to as the *clearance factor*, as shown in Figure 15-9.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

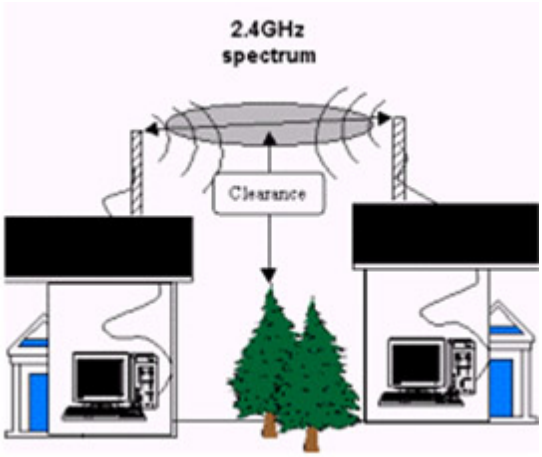

Figure 15-9: The distance of obstructive objects from the wireless signal is called the *clearance*
.

Distance — Although the actual distances over which antennas can send and receive data varies greatly based on different factors, you can achieve distances of many miles with outdoor antennas. For example, two Orinoco 24 dBi directional parabolic-grid antennas can theoretically achieve distances of up to 52 miles at an 11 Mbps transmission speed with a 180-meter clearance. Reduce that transmission rate to 1 Mbps and you can achieve distances of up to 149 miles with a 1200-meter clearance. Shorter distances are achieved with less expensive equipment, such as the Orinoco 14 dBi directional antenna which can achieve distances of up to 5.3 miles at 11 Mbps with a 13-meter clearance.

Cable factor — The distances that transmissions travel on the cables between the wireless cards and the antennas can be a factor in choosing the right antenna. The shorter the cables, the greater the distance and speed you will get on your antenna.

The power of an antenna is rated in terms of *gain*. Gain is measured in decibels, based on a *theoretic isotropic radiator* (or *dBi*). Higher gains offer opportunities for reaching greater distances at greater speeds. However, the ability of the antenna to focus that power (directional versus omnidirectional), greatly affects the speeds and distances that can be achieved.

Installing wireless Linux software

If you did a personal desktop or Everything installation of Red Hat Linux on your computer, the software packages you need to create your wireless LAN may already be installed. Drivers and modules needed to support PCMCIA cards and wireless cards should be in your system.

Besides the wireless drivers, the following software packages contain tools for configuring and working with your wireless LAN cards in Red Hat Linux:

kernel-pcmcia-cs — Contains commands and configuration files to support PCMCIA cards.

wireless-tools — Contains commands for setting extensions for your wireless LAN interface. Commands include `iwconfig` (for configuring your wireless interface) and `iwlist` (for listing wireless statistics).

Once you have established a wireless LAN interface, you can use a variety of Linux software to monitor and control access to that interface. You will need to install the appropriate software packages as well.

Configuring the wireless LAN

Before you begin testing the distances you can achieve with your wireless Linux LAN, I recommend that you configure wireless cards on two computers within a few feet of each other. Once the two computers are communicating, you can change wireless settings to tune the connection and begin experimenting with transmission distances.

The following sections describe the steps you need to take to set up a wireless LAN between two Linux systems. Although only two nodes are described, you can add more computers to your wireless LAN

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

once you know how. This procedure describes how to operate your wireless Linux LAN in two different modes:

Ad hoc — All the computers in your wireless LAN are gathered into a single virtual network made up of only one cell. A single cell means that you cannot roam among different groups of wireless nodes and continue your communication invisibly. To do that requires a managed network.

Managed — As I noted earlier, many wireless cards supported in Linux cannot operate as access points. A Linux wireless card, however, can operate as a node in a managed network. The wireless-configuration tools that come with Red Hat Linux let you identify the access point for Linux to use by indicating the access point's MAC address.

Choosing equipment

Start with two computers. (You can add more computers later, once you understand how to get your wireless interfaces working.) For this procedure, I used computers that had the following characteristics (you can use different computers and cards, if you like):

Computers — One computer was a laptop with an available PCMCIA slot; the other was a desktop system with only PCI slots.

Wireless cards — As I mentioned earlier, I purchased two Proxim (Lucent Technologies) Orinoco wireless LAN cards: one Gold Label and one Silver Label. For the desktop computer, I purchased a PCI adapter card because it had no PCMCIA slot. Both cards come with built-in antennas, so I had no need for additional antennas while setting up the two computers (in the same room) for wireless communication.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note The only difference between the Gold and Silver Label cards is that the Gold card offers support for more secure encryption, so there was no reason for me to choose two different typ

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Red Hat Linux — I installed Red Hat Linux on both machines, selecting a personal desktop install for the laptop and an Everything install for the desktop computer. (The Everything install was not strictly necessary, but I wanted to be able to use the GUI and various server features.)

Because I was using the desktop computer as a gateway to the Internet, that computer also had a wired Ethernet card that was connected to my DSL modem to provide a route to the Internet for any computers on my wired or wireless networks.

Inserting wireless cards

To physically install the wireless cards, follow the directions that come with the cards. For my laptop, I simply inserted one Orinoco card into a PCMCIA slot. For the desktop computer, I powered down, inserted the PCI adapter into a vacant slot, powered up, and inserted the other Orinoco PCMCIA card into the adapter.

Loading the modules

The `cardmgr` daemon monitors the PCMCIA slots on computers that have them. If a card is recognized and listed in the PCMCIA database when the card is inserted, the appropriate module is loaded. You should also hear two beeps indicating that the card has been recognized.

On my laptop, my Orinoco wireless card was recognized and its modules loaded. On my desktop computer (with the PCI adapter), the card was not recognized, so I had to do some extra configuration.

PCMCIA only

To see what modules are loaded after a card is inserted on a computer that has only PCMCIA slots, type the **lsmod** command (as root user). In my case, because the Orinoco card uses the orinoco and orinoco_cs modules (along with the hermes helper module), output from the `lsmod` command included the following lines:

```
# lsmod
Module          Size Used by     Not tainted
orinoco_cs       5640  1
orinoco         34368  0  [orinoco_cs]
hermes           5344  0  [orinoco_cs orinoco]
ds               8608  1  [orinoco_cs]
yenta_socket    12384  1
pcmcia_core     50752  0  [orinoco_cs ds yenta_socket]
```

You can see that the orinoco_cs module was loaded and that the referring modules included the orinoco module, the hermes module, the ds (PC Card Driver Services) module, and the pcmcia_core module. If you are using a different card, you may instead see one of the following modules: airo_cs, wavelan_cs, wvlan_cs, ray_cs, or netwave_cs.

PCMCIA with adapter card

If your computer has only ISA or PCI slots, you will need an adapter to use your PCMCIA wireless LAN card. Red Hat Linux detected my adapter card and added the following lines to the `/etc/sysconfig/pcmcia` file so that the PCI adapter card would be recognized and the PCMCIA service would start automatically at boot time:

```
PCMCIA=yes
PCIC=yenta_socket
```

The yenta_socket driver is a PCMCIA controller driver that includes the Yenta register specification. Yenta is used for CardBus bridges made by Cirrus Logic for a variety of manufacturers (Texas Instruments, IBM, Toshiba, and others). The Orinoco PCI adapter was detected as a device using the yenta_socket driver.
Figure 15-10 shows an example of an Orinoco Silver Card using a PCI adapter.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Figure 15-10: The Orinoco Silver wireless LAN card can be used with a PCI adapter (shown here).

Checking that the cards are working

If the modules have been loaded properly, the cardmgr should recognize each card and start up the Ethernet interface for it. To check that this has happened, restart the interface as follows:

/etc/init.d/pcmcia restart

Shutting down PCMCIA services: cardmgr modules.

Starting PCMCIA services: modules cardmgr.
You should hear a single beep when the card service stops, then two beeps when the adapter and wireless card are properly detected. Check the `/var/log/messages` file. You should see some messages at or near the end of this file, describing what happened when the PCMCIA interface was shut down and restarted. If the card is detected, you should see modules loaded successfully and a network interface started for the wireless card. Here are some examples:

Feb 9 17:26:33 toys kernel: Linux Kernel Card Services 3.1.22

Feb 9 17:26:33 toys kernel: options: [pci] [cardbus] [pm]

Feb 9 17:26:33 toys kernel: PCI: Found IRQ 5 for device 01:09.0

Feb 9 17:26:33 toys kernel: PCI: Sharing IRQ 5 with 00:1f.3

Feb 9 17:26:33 toys kernel: Yenta IRQ list 0000, PCI irq5

Feb 9 17:26:33 toys kernel: Socket status: 10000011

Feb 9 17:26:34 toys cardmgr[2571]: starting, version is 3.1.31

.

.

.

Feb 9 17:26:34 toys kernel: cs: IO port probe 0x0c00-0x0cff: clean.

Feb 9 17:26:34 toys kernel: cs: IO port probe 0x0100-0x04ff: excluding
0x400-0x47f 0x4d0-0x4d7

Feb 9 17:26:34 toys kernel: cs: IO port probe 0x0a00-0x0aff: clean.

Feb 9 17:26:34 toys kernel: cs: memory probe 0xa0000000-0xa0ffffff:
clean.

Feb 9 17:26:34 toys cardmgr[2571]: socket 0: Lucent Technologies
WaveLAN/IEEE Adapter

Feb 9 17:26:34 toys cardmgr[2571]: executing: 'modprobe hermes'

Feb 9 17:26:34 toys cardmgr[2571]: executing: 'modprobe orinoco'

Feb 9 17:26:34 toys cardmgr[2571]: executing: 'modprobe orinoco_cs'

Feb 9 17:26:34 toys cardmgr[2571]: executing: './network start eth1'

Feb 9 17:26:34 toys /etc/hotplug/net.agent: invoke ifup eth1

The preceding code shows that the kernel recognizes the PCI card (at IRQ 5). The cardmgr identifies

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

the Orinoco card as a WaveLAN/IEEE adapter in socket 0. The network script starts an Ethernet interface (eth1).
If the wireless LAN interface started properly, you should be able to see the new interface by using the `iwconfig` command. The following is an example of output from the `iwconfig` command:

```
eth1   IEEE 802.11-DS  ESSID:""  Nickname:"HERMES I"
       Mode:Managed  Frequency:2.457GHz  Access Point: 00:00:00:00:00:00
       Bit Rate:11Mb/s   Tx-Power=15 dBm   Sensitivity:1/3
       Retry limit:4   RTS thr:off   Fragment thr:off
       Encryption key:off
       Power Management:off
```

If your wireless LAN interface does not appear to be working, refer to the section "Troubleshooting a wireless LAN" later in this chapter. If the interface does seem to be working, you are ready to tune your wireless LAN card interface and set up TCP/IP to be able to use the interface.

Configuring the wireless interface

The Network Configuration window (`neat` command) can be used to configure wireless Ethernet card interfaces, as well as regular wired Ethernet cards. The following procedure describes how to configure a wireless Ethernet card using the Network Configuration window.

1. Start the Network Configuration. From the Red Hat menu, click System Settings → Network, or, as root user from a Terminal window, type **neat**. The Network Configuration window appears.
2. Click the New button. The Select Device Type window appears.
3. Click Wireless connection and Forward. The Select Wireless Device window appears.
4. Click your wireless card from the list of cards shown, and click Forward. The Configure Wireless Connection window appears, as shown in Figure 15-11.

Figure 15-11: Add a wireless interface using the Network Configuration window.

5. Add the following information and click Forward:

   Mode — Indicates the mode of operation for the wireless LAN card. Because I am setting up a wireless LAN consisting of only one cell (in other words, with no roaming to cells set up in other areas), I could set the mode to Ad hoc. Ad hoc mode allows the card to communicate directly with each of its peers. You can use Managed mode if you have multiple cells, requiring your card to communicate directly to an access point. You can also use Managed mode for a point-to-point network, such as when you use the wireless LAN to extend a network from one building to another.
   Network Name (SSID) — The network name (or Network ID) that identifies cells that are part of the same network. If you have a group of cells (which might include multiple nodes and repeaters among which a client could roam), this name can identify all of those cells

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

as falling under one virtual network. Choose any name you like and then use that name for all computers in your virtual network. (SSID stands for Service Set ID.)

Channel — Choose a channel between 1 and 14. You can begin with channel 1; if you get interference on that channel, try changing to other channels.

Transmit Rate — Choose the rate of transmission from the following rates: 11M, 5.5M, 2M, 1M, or Auto. Choosing Auto allows the interface to automatically ramp down to lower speeds as needed. Lower speeds allow the interface to transmit over greater distances and deal with noisy channels.

Key — You need the same encryption key for all wireless LAN cards that are communicating with each other. It is critical to get this value right. This key is used to encrypt all data transmitted and decrypt all data received on the wireless interface. You can enter the number (up to 10 digits) as XXXXXXXXXX or XXXX-XXXX-XX (where each X is a number), for example, 1234-5678-90.

A Configure Network Settings window appears.

6. You can enter the following information:

   Automatically obtain IP address settings with: If you want to get your IP address from a DHCP server, click this box and the rest of the information is obtained automatically. Otherwise, set the IP address statically using the other options.

   Host name: If you are using DHCP, you can optionally add a host name to identify this network interface. If none is entered here, the output from the `/bin/hostname` command is used.

   Statically set IP addresses: Click here to manually set your IP addresses.

   Address: If you selected static IP addresses, type the IP address of this computer into the Address box. This number must be unique on your wireless network.

   Subnet Mask: Enter the netmask to indicate what part of the IP address represents the network. (Netmask is described later in this chapter.)

   Default Gateway Address: If a computer on your wireless LAN is providing routing to the Internet or other network, type the IP address of the computer here.

7. Click Forward to see a listing of the information you just entered.
8. Click Apply to complete the new wireless network interface.
9. Click File → Save (on the main window) to save the interface.

This procedure creates an interface configuration file in your `/etc/sysconfig/network-scripts` directory. The name of the configuration file is `ifcfg-` followed by the interface name (such as `eth0`, `eth1`, and so on). So, if your wireless card is providing your only network interface, it would be called `ifcfg-eth0`.

Using any text editor, open the `ifcfg-eth`*?* file as root user. The following is an example of an `ifcfg-eth1` file:

```
# Please read /usr/share/doc/initscripts-*/sysconfig.txt
# for the documentation of these parameters.
USERCTL=no
PEERDNS=no
GATEWAY=10.0.0.1
TYPE=Wireless
DEVICE=eth1
HWADDR=00:02:2d:2e:8c:a8
BOOTPROTO=none
NETMASK=255.255.255.0
```

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

ONBOOT=no

IPADDR=10.0.1.1

NAME=

DOMAIN=

ESSID=

CHANNEL=1

MODE=Ad-Hoc

KEY=9900-0000-00

RATE=11Mb/s

NETWORK=10.0.1.0

BROADCAST=10.0.1.255

In this example, the wireless card's hardware (MAC) address is automatically set to `00:02:2d:2e:8c:a8`. (Your MAC address will be different.) The interface is not yet set to come up at boot time (`ONBOOT=no`). The interface device is `eth1` (which matches the interface filename `ifcfg-eth1`), because this particular computer has another Ethernet card on the `eth0` interface. The interface type is set to `Wireless`.

Other information in the file sets standard TCP/IP address information. The NETMASK is set to 255.255.255.0 and the IP address for the card is set to 10.0.1.1. The broadcast address is 10.0.1.255.

You can also set many options that are specific to your wireless network in this file. The following is a list of some additional options that you might want to set:

NWID — Identifies the name of this particular computer on the network. The computer's host name (determined from the `uname -n` command) is used by default if you don't set it with NWID.

FREQ — You can choose a particular frequency in which to transmit. No value is required, because selecting a channel implies a certain frequency. If you do enter a frequency, the value must be a number followed by a `k` (kilohertz), `M` (megahertz), or `G` (gigahertz). The default values for the channels you select range from `2.412G` (channel 1) to `2.484G` (channel 14), with other channels occurring at increments of .005G. The default is `2.422G`.

SENS —You can select the sensitivity level of the access point. SENS can be set to 1 (low density), 2 (medium density), 3 (high density). The default is 1. The sensitivity threshold has an impact on roaming.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Caution The encryption algorithm used with 802.11 networks is the Wired Equivalent Privacy (WEP) algorithm. Though using the encryption key is more secure than not using it, some experts

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Besides those options just shown, you can also pass any valid options to the `iwconfig` command (which actually interprets these values), by adding an IWCONFIG option to the configuration file. Display the `iwconfig` man page (`man iwconfig`) to see all wireless options. Also view the `/etc/sysconfig/network-scripts/ifup-wireless` script to see how the options you just added are processed.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Note On the computer that is acting as a gateway from your wireless network to the Internet, you need to turn on IP packet forwarding. Change the value of net.ipv4.ip_forward to 1 in /etc/sys

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

net.ipv4.ip_forward = 1

Repeat this procedure for each wireless Red Hat Linux computer on your LAN. At this point, your wireless network should be ready to go. Restart your network, as described in the following steps, to make sure that it is working.

Activating the wireless interfaces

To immediately activate the wireless interface you just configured, click on the Wireless entry on the Network Configuration window and click the Activate button. After a few seconds, the Status should appear as Active.

To have the interface start when you reboot your computer, click the wireless interface from the Network Configuration window and select Edit. From the Wireless Device Configuration window that appears, click the box next to "Activate device when computer starts."

If you want to explicitly enter a Network Name (SSID), click the Wireless Settings tab on the Wireless Device Configuration window. From there, select Specified, type the network name (any name you choose to match others on your wireless network), and click OK.

Be sure to save your changes on the Network Configuration window by clicking File → Save.

Checking your wireless connection

Your wireless LAN interface should be operating at this point. If another wireless computer is available on your wireless network, try communicating with it using the `ping` command and its IP address (as described in the " Can you reach another computer on the LAN?" section further in this chapter).
If you are not able to communicate with other wireless nodes or if transmission is slow, you may have more work to do. For example, if you see messages that say "Destination Host Unreachable," instead of the output shown earlier, refer to the section on "Troubleshooting a wireless LAN" for help. If you want to fine-tune your wireless interface, refer to the "Manually configuring wireless cards" section later in this chapter.

Wireless Security

The Wireless Ethernet Compatibility Alliance (WECA) has recommended changes in response to security concerns about wireless networks. They did this because, unlike wired networks, which can often be physically protected within a building, wireless networks often extend beyond physical boundaries that can be protected.

The Wireless Equivalent Privacy (WEP) standard adds encryption to the 802.11 wireless standard. WECA refers to WEP as its way of providing "walls" that make wireless Ethernet as secure as wired Ethernet. However, you need to implement WEP, as well as other security methods that would apply to any computer network, in order to make your wireless network secure. Here are WECA's suggestions:

- Change the default WEP encryption key on a regular basis (possibly weekly or even daily). This prevents casual drive-byhackers from reading your encrypted transmissions.
- Use password protection on your drives and folders.
- Change the default Network Name (SSID).
- Use session keys, if available in your product (session keys are not supported in current Linux wireless drivers).
- Use MAC address filtering (supported in a limited way in Linux).
- Use a VPN (Virtual Private Network) system, which can add another layer of encryption beyond that which is available on your wireless network.

For larger organizations requiring greater security, WECA suggests such features as firewalls and user-verification schemes (such as Kerberos). As I mentioned earlier in this chapter, features for

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

protecting from intrusions and restricting services are already built into Red Hat Linux. Refer to the descriptions of security tools in Chapters 14, 15, and 16 for methods of securing your network, its computers, and their services. In particular, you could consider adding a VPN such as CIPE (described in Chapter 16) to further secure all data sent on your wireless LAN.

Testing out distances

Although you may be thrilled to have a wireless LAN working between two computers, you will probably want these computers to be located some distance from each other to make the LAN useful. Getting your wireless LAN to work at the desired distances can be quite a challenge. See the section "Selecting antennas" earlier in this chapter for suggestions on selecting and using antennas to configure the type of wireless LAN you are interested in.

Setting wireless extensions

After the wireless module is loaded, you can change wireless extensions using the `iwconfig` command. The `iwconfig` command is the command that is actually used to set the options added to the `ifcfg` configuration script (for example, for the eth1 interface, the script would be `/etc/sysconfig/network-scripts/eth1`).

Some of the same options that you set when the module was loaded can be reset using the `iwconfig` command. The `iwconfig` command can be useful for testing different settings on an active wireless LAN. The syntax of the `iwconfig` command is as follows:

iwconfig *interface parameter value*

The *`interface`* is the name of the wireless interface you want to change, such as `eth1` or `wvlan0`. The *`parameter`* is the name of the option, and the *value* is replaced by its value. For example, to set your network name (ESSID) to Homelan, you could type the following as root user:

iwconfig eth0 essid "Homelan"

Table 15-2 contains a list of available options for the `iwconfig` command. Refer to the "Configuring the Wireless Interface" section for further details on these options.

Table 15-2: Options to the iwconfig Command

Option	Description
essid *name*	Indicates the network name.
ap *address*	Indicates that the access point is at a particular MAC address. For low-quality connections, the client driver may return to trying to automatically detect the access point. This setting is only useful in Managed mode.
channel #	Picks the channel number to operate on.
frag *frag_size*	Sets the fragmentation threshold for splitting up packets before they are transmitted.
freq 2.4*??***G**	Sets the frequency of the channel to communicate on.
key *xxxx-xxxx-xx*	Sets the key used for WEP encryption.
mode *option*	Sets the mode used for communications to Ad-hoc, Managed, Master, Repeater, Secondary, or Auto.
nick *name*	Sets the station name to define this particular computer.
rate *XX***M**	Defines the transmission rate to use.

ABC Amber CHM Converter Trial version, http://www.processtext.com/abcchm.html

Table 15-2: Options to the iwconfig Command

Option	Description
rts *number*	Sets the RTS/CTS threshold for packet transmission.
retry *number*	For cards that support MAC retransmissions, you can use this option to determine how many retries are made before the transmission fails. The value can be a number (indicating number of seconds allotted for retries), or a number followed by an `m` (for milliseconds) or `u` (for microseconds). Instead of a number, you can set a number of retries using the limit parameter. For example: `retry limit 100` indicates that the transmission can retry up to 100 times.
sens number	Sets the lowest possible sensitivity threshold for which the wireless interface will try to receive a packet. Raising this level can help block out interference from other wireless LANs that might weakly encroach on your transmission area.

The best place to add `iwconfig` options permanently in Red Hat Linux is the configuration file for your wireless interface in the `/etc/sysconfig/network-scripts` directory.
Options to `iwconfig` are added to the wireless interface file (such as `ifcfg-eth0` or `ifcfg-eth1`) using the `IWCONFIG` parameter. For example, to add an encryption-key value of 1234-1234-12 for your wireless LAN card, you could add the following line to your wireless-interface file:

```
IWCONFIG="key 1234-1234-12"
```

Team LiB

PREVIOUS NEXT