

THE EXPERT'S VOICE® IN OPEN SOURCE

Included CD
contains full version
of Ubuntu.

Beginning Ubuntu Linux

From Novice to Professional

The complete guide to Ubuntu—includes everything you need to begin taking advantage of this powerful Linux-based operating system!

Keir Thomas

Apress®

Beginning Ubuntu Linux

From Novice to Professional



Keir Thomas

Beginning Ubuntu Linux: From Novice to Professional

Copyright © 2006 by Keir Thomas

All rights reserved. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage or retrieval system, without the prior written permission of the copyright owner and the publisher.

ISBN-13 (pbk): 978-1-59059-627-2

ISBN-10 (pbk): 1-59059-627-7

Printed and bound in the United States of America 9 8 7 6 5 4 3 2 1

Trademarked names may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, we use the names only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

Lead Editor: Jason Gilmore

Technical Reviewer: John Hornbeck

Editorial Board: Steve Anglin, Dan Appleman, Ewan Buckingham, Gary Cornell, Jason Gilmore, Jonathan Hassell, James Huddleston, Chris Mills, Matthew Moodie, Dominic Shakeshaft, Jim Sumser, Matt Wade

Project Manager: Sofia Marchant

Copy Edit Manager: Nicole LeClerc

Copy Editor: Marilyn Smith

Assistant Production Director: Kari Brooks-Copony

Production Editor: Ellie Fountain

Compositor: Susan Glinert

Proofreader: Kim Burton

Indexer: Broccoli Information Management

Artist: Kinetic Publishing Services, LLC

Cover Designer: Kurt Krames

Manufacturing Director: Tom Debolski

Distributed to the book trade worldwide by Springer-Verlag New York, Inc., 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax 201-348-4505, e-mail orders-ny@springer-sbm.com, or visit <http://www.springeronline.com>.

For information on translations, please contact Apress directly at 2560 Ninth Street, Suite 219, Berkeley, CA 94710. Phone 510-549-5930, fax 510-549-5939, e-mail info@apress.com, or visit <http://www.apress.com>.

The information in this book is distributed on an “as is” basis, without warranty. Although every precaution has been taken in the preparation of this work, neither the author(s) nor Apress shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the information contained in this work.

Contents at a Glance

About the Author	xxiii
About the Technical Reviewer	xxv
Acknowledgments	xxvii
Introduction	xxix

PART 1 ■■■ Introducing the World of Linux

■ CHAPTER 1	Welcome!	3
■ CHAPTER 2	A History and Politics Lesson	11
■ CHAPTER 3	The Realities of Running Linux	19

PART 2 ■■■ Installing Ubuntu

■ CHAPTER 4	Preinstallation Steps	25
■ CHAPTER 5	Installing Ubuntu	33
■ CHAPTER 6	Solving Installation Problems	47

PART 3 ■■■ The No-Nonsense Getting Started Guide

■ CHAPTER 7	Booting Ubuntu for the First Time	63
■ CHAPTER 8	Getting Everything Up and Running	77
■ CHAPTER 9	How to Secure Your Computer	125
■ CHAPTER 10	Personalizing Ubuntu: Getting Everything Just Right	139
■ CHAPTER 11	Ubuntu Replacements for Windows Programs	161
■ CHAPTER 12	Managing Your Files	183

PART 4 ■■■ The Shell and Beyond

■ CHAPTER 13	Introducing the BASH Shell	203
■ CHAPTER 14	Understanding Linux Files and Users	219
■ CHAPTER 15	Working with Text Files	241
■ CHAPTER 16	Taking Control of the System	257
■ CHAPTER 17	Cool Shell Tricks	267

PART 5 ■■■ Multimedia

■ CHAPTER 18	Digital Music	279
■ CHAPTER 19	Movies and Multimedia	291
■ CHAPTER 20	Image Editing	301

PART 6 ■■■ Office Tasks

■ CHAPTER 21	Making the Move to OpenOffice.org	319
■ CHAPTER 22	OpenOffice.org Overview	329
■ CHAPTER 23	In Depth: Writer	343
■ CHAPTER 24	In Depth: Calc	355
■ CHAPTER 25	In Depth: Impress	363
■ CHAPTER 26	In Depth: Base	373
■ CHAPTER 27	In Depth: Evolution	383
■ CHAPTER 28	Running Microsoft Office Under Ubuntu	401

PART 7 ■■■ Keeping Your System Running

■ CHAPTER 29	Installing Software	413
■ CHAPTER 30	Managing Users	435
■ CHAPTER 31	Optimizing Your System	443
■ CHAPTER 32	Backing Up Data	463
■ CHAPTER 33	Scheduling Tasks	473
■ CHAPTER 34	Accessing Computers Remotely	479

■ APPENDIX A	Glossary of Linux Terms	493
■ APPENDIX B	BASH Command Index	515
■ APPENDIX C	Getting Further Help	529
■ APPENDIX D	Other Versions of Ubuntu	537
■ INDEX	541

Contents

About the Author	xxiii
About the Technical Reviewer	xxv
Acknowledgments	xxvii
Introduction	xxix

PART 1 ■■■ Introducing the World of Linux

CHAPTER 1	Welcome!	3
	What Is Linux?	3
	The Age of Linux	5
	The Problems with Windows	6
	The Benefits of Linux	8
	Crash-Free	8
	Security	8
	Free and Shareable	8
	Compatible with Older Hardware	8
	The Linux Community	9
	Summary	10
CHAPTER 2	A History and Politics Lesson	11
	In the Beginning	11
	The GNU Project	13
	Proprietary Software and the GPL	14
	GNU and Linux Together	15
	Different Flavors of Linux	16
	Linux Today	17
	Modern Linux Development	17
	Summary	18

CHAPTER 3	The Realities of Running Linux	19
	Learning to Use Linux	19
	Who Uses Linux?	20
	Getting Hold of Linux	21
	Using Ubuntu	22
	Summary	22

PART 2 ■■■ Installing Ubuntu

CHAPTER 4	Preinstallation Steps	25
	Understanding Partitioning	25
	Freeing Up Space	27
	Reclaiming Space	28
	Removing Windows	28
	Using Another Hard Disk	28
	Backing Up Your Data	29
	Backing Up E-mail Files	30
	Making Notes	31
	Summary	32
CHAPTER 5	Installing Ubuntu	33
	An Overview of the Installation Process	33
	Step-by-Step Guide	34
	Step 1: Boot from the CD-ROM	34
	Step 2: Select Your Language	36
	Step 3: Confirm Your Keyboard Layout	36
	Step 4: Wait During Hardware and Networking Autodiscovery	37
	Step 5: Set a Hostname	37
	Step 6: Partition the Disk	38
	Step 7: Wait During Automated Package Copying	40
	Step 8: Set the Time Zone	41
	Step 9: Set Your Real Name, Username, and Password	41
	Step 10: Install the GRUB Boot Loader	42
	Step 11: Reboot for the Second Installation Phase	43
	Step 12: Wait During Package Installation	44
	Step 13: Configure Your Monitor	44
	Step 14: Boot for the First Time	45
	Summary	46

CHAPTER 6	Solving Installation Problems	47
	Preinstallation Problems	47
	Installation Problems	49
	Postinstallation Problems	51
	Graphical Problems	53
	Graphics Card Configuration	54
	Keyboard and Mouse Configuration	55
	Monitor Configuration	57
	Summary	59

PART 3 ■■■ The No-Nonsense Getting Started Guide

CHAPTER 7	Booting Ubuntu for the First Time	63
	Starting Up	63
	Logging In	64
	Exploring the Desktop	66
	First Impressions	66
	Desktop Elements	68
	Quick Desktop Guides	69
	Running Programs	71
	Working with Virtual Desktops	72
	Using the Mouse	73
	Cutting and Pasting Text	74
	Summary	75
CHAPTER 8	Getting Everything Up and Running	77
	Ubuntu Hardware Support	77
	Proprietary vs. Open Source Drivers	78
	Viewing Your Hardware	79
	Getting Online	81
	Using an Ethernet Card	81
	Joining a Wireless Network	84
	Configuring Dial-Up Internet Access	90
	Using a USB-Based DSL Modem	92
	Working with a Proxy Server	98

Setting Up Online Software Repositories	99
Configuring E-Mail and Instant Messaging	103
Configuring E-Mail Access	103
Setting Up Instant Messaging	105
Adding a Printer	107
Configuring a Local Printer	107
Configuring a Network Printer	108
Configuring a Windows/SMB Shared Printer	109
Setting Up Digital Imaging Devices	111
Transferring Photos from a Digital Camera	111
Configuring a Scanner	114
Using a USB Memory Stick	116
Installing 3D Graphics Card Drivers	116
Installing an ATI Driver	117
Installing a Nvidia Driver	117
Testing 3D Capabilities	118
Configuring Bluetooth	118
Pairing Bluetooth Devices	119
Transferring Files Between Bluetooth Devices	120
Using a Bluetooth Keyboard or Mouse	122
Configuring Sound Cards	123
Summary	124

CHAPTER 9 How to Secure Your Computer

Windows Security vs. Linux Security	125
Root and Ordinary Users	126
Common-Sense Security	128
Online Updates	130
The Ubuntu Firewall	131
Installing Firestarter	131
Configuring Firestarter	132
Summary	137

CHAPTER 10 Personalizing Ubuntu: Getting Everything Just Right

Changing the Look and Feel	139
Altering the Theme	139
Changing the Wallpaper	143
Setting Font Preferences	144
Configuring Input Devices	146

Creating Keyboard Shortcuts.....	148
Personalizing Login Options	149
Changing Your Login Picture.....	151
Adding and Removing Desktop Items	152
Adding a Shortcut.....	152
Adding and Removing Menus	154
Personalizing the Panels	156
Working with Applets	156
Setting Power-Saving Features	157
Controlling the Monitor Shutoff.....	157
Spinning Down the Hard Disk	158
Summary	159
CHAPTER 11 Ubuntu Replacements for Windows Programs	161
Available Software	161
A Quick Start with Common Linux Programs	163
Word Processing: OpenOffice.org Writer	163
Spreadsheet: OpenOffice.org Calc	165
Presentations: OpenOffice.org Impress	166
Database: OpenOffice.org Base	167
E-Mail/Personal Information Manager: Evolution	168
Web Browser: Firefox	169
Audio Playback: Rhythmbox and CD Player.....	171
Movie Playback: Totem Movie Player.....	172
CD/DVD Burning: Nautilus and Serpentine	173
Photo Editing: The GIMP.....	174
Other Handy Applications.....	175
Summary	181
CHAPTER 12 Managing Your Files	183
Understanding File System Concepts	183
Drive References.....	183
Case Sensitivity.....	184
File Access and Storage.....	184
Using Nautilus	186
File and Folder Icons	188
Special Nautilus Windows	189
Launching Files and Running Programs	190

Accessing Windows Files	192
Working with Files in Windows Partitions	192
Accessing Networked Files	195
Accessing Removable Storage Devices	196
Ejecting Media from Drives	197
Formatting Floppies	198
Summary	199

PART 4 ■■■ The Shell and Beyond

■ CHAPTER 13 Introducing the BASH Shell	203
What Is the BASH Shell?	203
Why Bother with the Shell?	205
When Should You Use the Shell?	206
Getting Started with the Shell	207
Running Programs	208
Getting Help	209
Running the Shell via a Virtual Console	210
Working with Files	211
Listing Files	213
Copying Files	214
Moving Files	215
Deleting Files	216
Changing and Creating Directories	217
Summary	217
■ CHAPTER 14 Understanding Linux Files and Users	219
Real Files and Virtual Files	219
Users and File Permissions	221
Viewing Permissions	223
Altering Permissions	227
The File System Explained	228
Mounting	232
Using the mount Command	232
Mounting a Drive Manually	233
Removing a Mounted System	234

File Searches	235
Using the find Command	235
Using the locate Command	236
Using the whereis Command	237
File Size and Free Space	237
Viewing File Sizes	237
Finding Out the Amount of Free Space	238
Summary	239
CHAPTER 15 Working with Text Files	241
Viewing Text Files	241
Using the cat Command	241
Using the less Command	241
Using the head and tail Commands	242
Using a Command-Line Text Editor	244
Understanding vi Modes	244
Using vi to Edit a File	248
Creating a New Text File Using vi	250
Searching Through Files	251
Using grep to Find Text	251
Using Regular Expressions	252
Comparing Text Files	254
Summary	255
CHAPTER 16 Taking Control of the System	257
Viewing Processes	257
Controlling Processes	260
Killing Processes	260
Controlling Zombie Processes	262
Using Other Commands to Control Processes	263
Controlling Jobs	264
Summary	265
CHAPTER 17 Cool Shell Tricks	267
Using Autocompletion	267
Autocompletion with Files and Paths	268
Viewing Available Options	268

Using Keyboard Shortcuts	269
Shortcuts for Working in BASH	269
Shortcuts for System Control	270
Using the Command History	271
Piping and Directing Output	272
Piping the Output of Commands	272
Redirecting Output	274
Summary	276

PART 5 ■■■ Multimedia

CHAPTER 18 Digital Music	279
Understanding the Legalities of Playback	279
Playing Music Files	280
Installing Codecs	280
Using Rhythmbox	282
Listening to Audio CDs	283
Tuning in to Online Radio Stations	284
Configuring XMMS	284
Using Streamtuner	284
Ripping Music from CDs	285
Choosing a Format	285
Adding MP3 Support to Sound Juicer	286
Ripping Tracks	287
Creating Your Own CDs	288
Summary	289
CHAPTER 19 Movies and Multimedia	291
Watching Videos	291
Installing Codecs	292
Installing RealPlayer	293
Playing Back Video	293
Watching DVDs	296
Playing Flash Animations	297
Watching TV	299
Checking for Video Input	299
Installing Tvtime	300
Summary	300

CHAPTER 20	Image Editing	301
	Getting Pictures onto Your PC	301
	Introducing The GIMP	302
	Editing Images with The GIMP	306
	Making Color Corrections	308
	Cropping and Cloning	309
	Sharpening	310
	Applying Filters	311
	Summary	315

PART 6 ■■■ Office Tasks

CHAPTER 21	Making the Move to OpenOffice.org	319
	Office Similarities	319
	OpenOffice.Org Key Features	320
	File Compatibility	321
	The Right Fonts	323
	Copying Windows Fonts	323
	Installing TrueType Core Fonts	324
	Summary	327
CHAPTER 22	OpenOffice.org Overview	329
	Introducing the Interface	329
	Customizing the Interface	330
	Adding Functions to Toolbars	330
	Adding a New Toolbar	333
	Customizing Menus	333
	Configuring OpenOffice.org Options	334
	Using OpenOffice.org Core Functions	339
	Using Wizards	339
	Getting Help	339
	Inserting Objects with Object Linking and Embedding	339
	Creating Macros	340
	Saving Files	341
	Summary	342

CHAPTER 23	In Depth: Writer	343
	Formatting Text	343
	The Formatting Toolbar	343
	Context Menu	344
	The Style and Formatting Palette	345
	Spell-Checking	345
	Inserting Pictures	346
	Working with Tables	347
	Mail Merging	349
	Adding Headers and Footers	352
	Summary	353
CHAPTER 24	In Depth: Calc	355
	Entering and Formatting Data	355
	Deleting and Inserting Data and Cells	356
	Working with Formulas	356
	Using the Function Wizard	357
	Summing Figures	358
	Sorting Data	359
	Creating Charts	359
	Using Filters	361
	Summary	362
CHAPTER 25	In Depth: Impress	363
	Creating a Quick Presentation	363
	Working in Impress	365
	Animating Slides	366
	Applying Fontwork	367
	Using 3D Effects	369
	Exporting a Presentation As a Flash File	370
	Summary	371

CHAPTER 26 In Depth: Base	373
Getting Started with Base	373
Creating a Database	375
Adding a Table	375
Creating a Form	377
Adding Controls to the Form Manually	378
Editing the Database Table	379
Using the Database	380
Summary	381
CHAPTER 27 In Depth: Evolution	383
Evolution Modes	383
Basic E-Mail Tasks	385
Sending and Receiving E-Mail	385
Reading E-Mail	386
Deleting Messages	386
Flagging Messages	387
Composing a Message	388
Creating an E-Mail Signature	389
Advanced E-Mail Tasks	390
Creating New Folders	390
Dealing with Junk E-Mail	391
Sorting and Filtering Messages	392
Creating Search Folders	394
Contacts	394
Adding or Editing Contact Information	395
Creating a Contact List	396
Calendars	396
Specifying Appointment Types	396
Adding or Editing a Diary Entry	397
Tasks	398
Summary	399

CHAPTER 28	Running Microsoft Office Under Ubuntu	401
	Installing and Configuring Wine	401
	Downloading Wine	401
	Configuring Wine	403
	Showing Hidden Files on CDs	406
	Installing Microsoft Office	406
	Running Office Components	407
	Installing Other Windows Programs	409
	Summary	410

PART 7 ■■■ Keeping Your System Running

CHAPTER 29	Installing Software	413
	Software Installation Basics	413
	Formats of Linux Installation Files	414
	Packages Explained	415
	Dependency Management	416
	APT and Package Repositories	417
	The Synaptic Package Manager in Depth	420
	Searching for Software	420
	Installing Software	422
	Removing Software	424
	Package Management from the Command Prompt	424
	Using dpkg	424
	Using the APT Tools	426
	Installing from Source	430
	Installing the Compiler Tools	430
	Unpacking the Tarball and Solving Dependencies	431
	Compiling	432
	Summary	434
CHAPTER 30	Managing Users	435
	Understanding User and Group Accounts	435
	Users and Groups	435
	Root User	436
	UIDs and GIDs	437

Adding and Deleting Users and Groups	437
Adding and Deleting Users via the GUI	437
Creating and Deleting Groups via the GUI	439
Adding and Deleting Users and Groups at the Command Line . . .	439
Adding and Changing Passwords	441
Summary	442
CHAPTER 31 Optimizing Your System	443
Speeding Up Booting	443
Understanding Bootup	443
Reducing the Boot Menu Delay	445
Optimizing Initialization and Run-Level Services	446
Optimizing Hard Disk Settings	452
Benchmarking Your Hard Disk	452
Changing Hard Disk Settings	453
Making Disk Optimizations Permanent	455
Prelinking	456
Configuring Prelinking	457
Deactivating Prelinking	457
Freeing Disk Space	458
Emptying the /tmp Folder	458
Emptying the Cache of Package Files	459
Removing Unused Software	460
Adding Another Disk Drive	460
Partitioning the Disk	460
Configuring Ubuntu to Use the Drive	461
Summary	461
CHAPTER 32 Backing Up Data	463
What Data Should You Back Up?	463
Using SBackup	464
Backing Up Data via SBackup	465
Restoring Data via SBackup	468
Making Backups from the Command Line	468
Creating tar Files	469
Compressing tar Archives	469
Extracting Files from a tar Archive	470
Viewing tar Archive Information	470
Saving the File to a CD-R/RW	471
Summary	471

CHAPTER 33	Scheduling Tasks	473
	Scheduling with crontab	473
	Creating a Scheduled Task	474
	Editing crontab with the GUI	475
	Scheduling with anacron	476
	Using at to Schedule One-Off Tasks	477
	Summary	478
CHAPTER 34	Accessing Computers Remotely	479
	Using Secure Shell	479
	Logging In to a Remote Computer	480
	Transferring Files Between Remote Computers	483
	Accessing GUI Applications Remotely	485
	Running X Applications on a Remote Computer	485
	Running the GNOME Desktop Across a Remote Connection	486
	Accessing Ubuntu via Remote Desktop	488
	Connecting to Remote Windows Computers	489
	Connecting to Windows XP Professional, 2000, and NT	489
	Connecting to Other Windows Computers	491
	Summary	492
APPENDIX A	Glossary of Linux Terms	493
APPENDIX B	BASH Command Index	515
APPENDIX C	Getting Further Help	529
	Read the Manual!	529
	Man Pages	529
	Info Pages	532
	README Files and Other Documentation	533
	Getting Help Online	533
	Forums	533
	Mailing Lists	534
	Other Official Sites	535
	Third-Party Sites	535

APPENDIX D Other Versions of Ubuntu	537
Kubuntu	537
Edubuntu	538
Ubuntu Server	538
Xubuntu	539
Other Ubuntu Projects	540
INDEX	541

About the Author



■ **KEIR THOMAS** has been writing about computers, operating systems, and software for a decade. He has edited several best-selling computer magazines, including *LinuxUser & Developer*, *PC Utilities*, and *PC Extreme*, and worked as part of the editorial staff on a range of other titles. He was formerly Technical Group Editor at Live Publishing. Throughout Keir's career, his aim has been to explain advanced and confusing technology in ways that the average person can understand. Keir works as a freelance editor and writer. He lives on the side of a mountain in England, and his pastimes include hiking and playing musical instruments. *Beginning Ubuntu Linux* is Keir's second title for Apress. His first was *Beginning SUSE Linux*.

About the Technical Reviewer



■ **JOHN HORNBECK** has been involved in free software since 1999 and involved with Ubuntu since before the first release. He was a member of the original Ubuntu Documentation Team and was responsible for writing the documentation that shipped with the first Ubuntu release. John was also one of the first Ubuntu developers not directly employed by Canonical Ltd.

John currently runs his own software company, Oklahoma Computer, Software and Robotics, and attends university classes. When not writing code, he spends time with his wife and two children.

Acknowledgments

Books like the one you're holding now take an enormous amount of work by a lot of people to come to fruition. To this end, I'd like to acknowledge the help of the following individuals in the production of this book: Chris Mills, Emily K. Wolman, Frank Pohlmann, Jason Gilmore, John Hornbeck, Sofia Marchant, Marilyn Smith, Ellie Fountain, and the many other people at Apress whose work behind the scenes made *Beginning Ubuntu Linux* possible.

In addition, I'd like to thank the Ubuntu community at large, who often helped me figure out some aspects of Ubuntu, particularly when it came to hardware configuration. Thanks must also go to Mark Shuttleworth and the Ubuntu developers for creating one of the most exciting and innovative Linux distributions I've had the pleasure of using.

Introduction

Linux has come a long way in a short time. Computing itself is still relatively young by any standard; if the era of modern computing started with the invention of the microchip, it's still less than 50 years old. But Linux is a youngster compared even to this; it has been around for only 15 of those years.

In that brief time span, a student's personal project has grown to where it now runs many computers throughout the world. It has rampaged through the computing industry, offering an alternative to commercial solutions such as those offered by Microsoft, and toppling long-held beliefs about the way things should be done. This is all by virtue of the fact that Linux is simply better than every other choice out there. Many argue that it's more secure and faster than other operating systems. But here's the kicker—Linux is free of charge. Yes, that's right. It doesn't have to cost a penny. It is one of the computing industry's best-kept secrets.

I was bitten by the Linux bug in the mid-1990s. I was introduced to it by a friend who sold it to me as a kind of alternative to DOS. At that time, I tapped a few commands at the prompt and was greeted by error messages. I must admit that I was put off. But shortly afterwards, I revisited Linux and quickly became hooked.

Yet getting used to Linux wasn't easy. I read as many books as I could, but they weren't very helpful to me. They were usually overly complicated or simply irrelevant. To start off, I didn't want to know how to create a program that could parse text files. I just wanted to know how to copy and delete files. I didn't want to set up a web server. I just wanted to know how to play my MP3 tracks and browse the Web.

This book is my answer to the need for a fundamental, authoritative, and down-to-earth guide to Linux, done in the context of one of the most popular flavors of Linux in existence today. It's a book that is desperately needed in our modern world, especially as Linux becomes more and more popular and enters homes and workplaces.

Beginning Ubuntu Linux purely and simply focuses on what you need to know to use Linux. It's concise and to the point, aiming to re-create under Linux all the stuff you used to do under Windows. But don't think that this means *Beginning Ubuntu Linux* cuts corners. Wherever justified, this book spends time examining the topics you need to know in order to gain a complete and comprehensive understanding. For example, you'll find a hefty chapter looking at the command-line prompt—arguably the heart of Linux and the element that gives Linux most of its power. There's also an entire chapter discussing (and illustrating) how to initially install Ubuntu on your computer. *Beginning Ubuntu Linux* really is a complete guide.

About Ubuntu

Linux applies an alternative philosophy to computing that revolves around the sharing of not only software, but also knowledge. To use Linux is to become part of a huge global community of people who have caught on to a phenomenon that is changing the world.

Ubuntu (www.ubuntulinux.org) is the natural continuation of these goals. It's a project founded by entrepreneur businessman Mark Shuttleworth with the intention of bringing a freely available, high-quality operating system to the world. To this end, Shuttleworth invested \$10 million of his own money in order to guarantee that this will be the case for many years to come.

The fundamental concept is that Ubuntu is available for use by anyone in the world, no matter who they are or where they are. As such, many different languages are supported, and the operating system can also be accessed by those with disabilities, such as partial sight or hearing. Ubuntu might just as easily be found on a Wall Street banker's laptop as on a battered, old computer in a Brazilian favela.

Ubuntu is built around one of the most established versions of Linux: Debian (www.debian.org). The Debian Project was started back in 1993, shortly after the very first version of the Linux software was released, and has risen to become one of the pioneering varieties of Linux.

Ubuntu and Debian Linux both share common goals and are closely allied, but Ubuntu focuses largely on the desktop. It provides a powerful office suite by default, for example, as well as some excellent pieces of Internet software.

It's also very easy to use. Ubuntu works "straight out of the box." As soon as it's installed, you should be ready to start using it without any further work. In addition, tasks such as updating your software are as easy under Ubuntu as they are under Windows.

Above all, however, Ubuntu is designed to be shared. You can take the CD-ROM included with this book and install Ubuntu on as many machines as you want. You can also copy it as many times as you want and give those copies to your friends. This isn't some kind of trick—Ubuntu isn't a trial version that will shut up shop in a month. You will never find yourself having to pay a fee further down the line, even if you want to install additional software. Ubuntu will always be free of charge.

What You'll Find in This Book

Beginning Ubuntu Linux is split into seven parts, each of which contains chapters about a certain aspect of Ubuntu use. These parts can be read in sequence, or you can dip in and out of them at will. Whenever a technical term is mentioned, a reference is made to the chapter where that term is explained.

Part 1 examines the history and philosophy behind the Linux operating system. I aim to answer many of the common questions about Linux. Such knowledge is considered to be as important, if not more so, than understanding the technical details on how Linux works. But while these chapters should be read sooner rather than later, they don't contain any technical information that you absolutely require to get started with Ubuntu.

Part 2 covers installing Ubuntu on your computer. An illustrated guide is provided, and all installation choices are explained in depth. Additionally, you'll find a problem-solving chapter to help in case anything goes wrong.

Part 3 focuses on getting started with Ubuntu. It covers setting up the Linux system so that it's ready to use. One chapter is dedicated to setting up common hardware devices, such as printers and modems, and another explains how you can secure your system. Other chapters in this part explore the desktop, explaining what you need to know to begin using Ubuntu on a daily basis.

In Part 4, we take a look at how the underlying technology behind Linux functions. You'll be introduced to the command-line prompt and learn how the file system works. It's in these chapters that you'll really master controlling Linux!

Part 5 covers multimedia functions built into Ubuntu, which let you watch movies and play back music. We also take a look at the image-editing software built into Ubuntu.

Part 6 moves on to explain how typical office tasks can be accomplished under Ubuntu. We investigate OpenOffice.org, the complete office suite built into Ubuntu. After an introduction to OpenOffice.org, separate chapters explore its word processor, spreadsheet, database, and presentation package. You'll also learn how to use Evolution, Ubuntu's e-mail program, and how to run Microsoft Office under Ubuntu.

Part 7 carries on from Part 4, taking an even more in-depth look at the underlying technology behind Ubuntu. This time, the emphasis is on giving you the skills you need to keep your system running smoothly. You'll learn how to install software, manage users, optimize your system, back up essential data, schedule tasks, and access computers remotely.

Finally, at the back of the book, you'll find four appendixes. The first is a glossary of Linux terms used not only in this book, but also in the Linux and Unix world. The second appendix is a quick reference to commands typically used at the command-line prompt under Linux. The third appendix explains how to get further help when using Ubuntu. The final appendix briefly describes some other versions of Ubuntu that differ from the one accompanying this book.

Conventions Used in This Book

The goal when writing *Beginning Ubuntu Linux* was to make it as readable as possible while providing the facility for readers to learn at their own pace.

Throughout the book, you'll find various types of notes and sidebars complementing the regular text. These are designed to provide handy information to help further your knowledge. They also make reading the book a bit easier.

■ **Note** A note is designed to provide an important piece of information that you should know and that will help your understanding of the topic being discussed.

■ **Tip** A tip is something that will help when you need to perform the task being described. Alternatively, it might be something that can make your life easier when using Ubuntu.

■ **Caution** A caution is something you should certainly pay attention to because it warns of a hidden danger or particular caveat that applies to the topic being discussed.

In the sidebars, I take a moment to explain something that you should know but isn't vital to an understanding of the main topic being discussed. You don't need to read the sidebars there and then; you can return to them later if you wish.

Support from Apress

Finally, I'd like to mention the online help offered by Apress, the publisher of this book. When you visit www.apress.com, you can click a link at the top of the page to visit the forums. Here, you'll be able to make comments about this book and find help with any issues it raises.

PART 1



Introducing the World of Linux



Welcome!

If you're an avid computer user, there's a good chance that you've heard of Linux. You might have read about it, or perhaps you've seen TV ads that refer to it.

One of the odd things about Linux is that the more you learn about it, the more questions you have. For instance, it's generally thought that Linux is free of charge, but this then raises the question of how, in our modern world, something like an entire computer operating system can cost nothing. Who pays the programmers?

Over the following introductory chapters, I'm going to try to answer some of these questions. In this chapter, I'll explain what Linux is and its benefits compared to Windows.

What Is Linux?

There are two ways of looking at a PC. The first is to see it as a magical box, which lets you do cool stuff like browse the Internet or play games. Seen in this way, it's like a VCR—put in a tape, press a button, and a picture appears on your TV. On your PC, you click the Internet Explorer icon, type a web address, and a web site somehow appears. The astonishing technical complexity behind these simple procedures isn't important to most people.

The other way of looking at a PC is as a collection of components that are made by various manufacturers. You might be familiar with this way of thinking if you're ever tried to upgrade your PC's hardware. In that case, you'll know that your PC consists of a CPU, a hard disk, a graphics card, and so on. You can swap any of these out to put in newer and better components that upgrade your PC's performance or allow more data storage.

What almost no one realizes is that the operating system is just another component of your PC. It, too, can be swapped out for a better replacement. Windows doesn't come free of charge, and Microsoft isn't performing a public service by providing it. Around \$50 to \$100 of the price you pay for a PC goes straight into Microsoft's pocket. Bearing in mind that hundreds of millions of PCs are made each year, it's not hard to see why Microsoft is one of the world's richest corporations.

It would be difficult to question this state of affairs if Microsoft gave us our money's worth. But it often falls far short. Its products are full of serious security holes, which at best inconvenience us and at worst make us lose data. And that's before we consider the instability of Microsoft products—hardly a day goes by without something unexpected happening. One of the first things people are taught when attending Windows training is how to use the Ctrl+Alt+Delete keyboard combination, which helps you recover after a crash!

Microsoft became rich, and maintains its wealth, by a virtual monopoly over PC manufacturers. While the intelligent computer buyer can choose between components to put together a better PC—deciding between an AMD or Intel processor, for example—you usually have little choice but to buy Windows with a new PC. Try it now. Phone your favorite big-name computer retailer. Say that you want a PC but you *don't* want Windows installed. Then listen as the salesperson on the other end of the phone struggles to understand.

Note Some PC manufacturers actually will sell you a PC without Windows installed on it. All you have to do is ask, although you might need to speak to a senior salesperson to get through to someone who understands your request. Smaller local companies, in particular, will be more than willing to sell you a PC without Windows. Some larger multinational companies, such as Hewlett-Packard, sell workstations with Linux preinstalled instead of Windows. However, these computers are usually aimed at businesses rather than home users.

Wouldn't it be terrific if you could get rid of Windows? Would you like to finally say goodbye to all those security holes and not have to worry about virus infections anymore, yet not lose out on any features or need to make sacrifices or compromises?

There is an alternative. Welcome to the world of Linux.

Linux is an operating system, which is to say that it's a bit like Windows. It's the core software that runs your computer and lets you do stuff on it. By the strictest definition of the term, an *operating system* is the fundamental software that's needed to make your PC work. Without an operating system installed on your PC, it would merely be an expensive doorstop. When you turned it on, it would beep in annoyance—its way of telling you that it can't do much without a whole set of programs to tell it what to do next.

An operating system allows your PC's hardware to communicate with the software you run on it. It's hundreds of programs, system libraries, drivers, and more, all tightly integrated into a whole. In addition, an operating system lets programs talk to other programs and, of course, communicate with you, the user. In other words, the operating system runs everything and allows everything to work.

Note Some companies and individuals, including Microsoft, define an operating system as much more than this fundamental software. They add in the basic tools you run on an operating system, such as web browsers and file management programs.

Linux consists of a central set of programs that run the PC on a low level, referred to as the *kernel*, and hundreds (if not thousands) of additional programs provided by other people and various companies. Technically speaking, the word *Linux* refers explicitly just to the core kernel program. However, most people generally refer to the entire bundle of programs that make up the operating system as *Linux*.

GNU/LINUX

Although most of us refer to Linux as a complete operating system, the title “Linux” hides a lot of confusing but rather important details. Technically speaking, the word *Linux* refers merely to the kernel file: the central set of programs that lie at the heart of the operating system. Everything else that comes with a typical version of Linux, such as programs to display graphics on the screen or let the user input data, is supplied by other people, organizations, or companies. The Linux operating system is the combination of many disparate projects. (I’ll explain how this works in the next chapter.)

The GNU organization, in particular, supplies a lot of vital programs and also system library files, without which Linux wouldn’t run. These programs and files were vital to the acceptance of Linux as an operating system in its early days. Because of this, and the fact that Linux completed a long-running goal of the GNU project to create a Unix-like operating system, some people choose to refer to Linux as GNU/Linux.

A fierce debate rages over the correct way to refer to the Linux operating system and whether the GNU prefix should be used. For what it’s worth, an equally fierce debate rages over how we should define an operating system. It can all get very confusing. It’s also very easy to accidentally offend someone by not using the correct terminology!

It’s not the purpose of this book to get involved in this debate. Suffice it to say that I acknowledge the vital input of the GNU project into the operating system many people refer to simply as Linux, as well as that of other vital projects. However, readers should note that when I refer to *Linux* throughout this book, I mean the entire operating system. If I intend to refer simply to the kernel programs, I will make that clear.

The Age of Linux

At the time of writing this book, Linux is a little over 14 years old. It has gone from a hobbyist project maintained by just one man to a professional and corporate-sponsored solution for virtually every level of computer user.

Linux has also gone from being a server operating system, designed for central computers that hand out files and other computer resources to other computers, to becoming a full-fledged graphical desktop operating system like Windows. In fact, it’s gone even further. Today, it’s very likely that you’ll find Linux running your digital video recorder and other computerized household gadgets.

Getting technical for a moment, Linux is a 32-bit and 64-bit, multitasking, multiuser operating system. This is a complicated way of saying that it’s pretty darn powerful. Linux is as capable of running supercomputers as it is of running a desktop PC. Linux builds on the foundation laid by Unix, which itself was based on Multics, which was one of the first modern computer operating systems. It’s not an exaggeration to say that Linux can trace its family tree all the way back to the pioneering days of computing.

CORRECT PRONUNCIATION

What most people refer to as the Linux operating system takes its name from the kernel program, one of its most important system components. This, in turn, was named after its creator, Linus Torvalds.

The name Linus is commonly pronounced “Lie-nus” in many English-speaking countries, but Torvalds speaks Swedish. He pronounces his name “Leen-us” (imagine this spoken with a gentle Scandinavian lilt, and you’ve got it about right).

Because of this, he pronounces Linux as “Lin-ux” and most people copy this pronunciation. You can hear this spoken by Torvalds himself by visiting www.paul.sladen.org/pronunciation/.

Some people refer to the Linux operating system by its full title of GNU/Linux. In this case, GNU is pronounced as in the name of the animal, with a hard G: “G-noo.” The full pronunciation is therefore “G-noo Lin-ux.”

Finally, the CD that comes with this book contains a version of Linux called Ubuntu. Ubuntu is an African word that in its country of origin is pronounced in three separate short syllables: “oo-bu-ntoo.” However, most western English speakers pronounce the word “oo-bunt-oo,” which is perfectly acceptable!

The Problems with Windows

The world’s most popular operating system is Windows, which is made by the Microsoft Corporation. Linux has no links with Windows at all. Microsoft doesn’t contribute anything to Linux and, in fact, is rather hostile toward it because it threatens Microsoft’s market dominance. This means that installing Linux can give you an entirely Microsoft-free PC. How enticing does *that* sound?

Windows is used on 91% of the world’s desktop computers. In other words, it must be doing a good job for it to be so popular, right?

Let’s face facts. Windows is not without problems, and that’s putting it mildly. It’s stunningly insecure, and virtually every day a new security hole is uncovered. This leads to the creation of worms by malevolent programmers. *Worms* are small programs that exploit security holes within operating systems, leaping from computer to computer and spreading like wildfire via the Internet. Examples include Sasser (as well as its variations), which causes your computer to crash and shut down as soon as you go online.

Then there are the viruses—hundreds and hundreds of them. This has led to an entire industry that creates antivirus programs, which are additional pieces of software. Antivirus software is vital if you want to use Windows without losing data or running the risk of your files being stolen!

Some argue that Windows is hit by so many viruses merely because it’s so popular. But consider that many of these viruses are simple programs that merely take advantage of security holes in Windows. For example, one particular virus took advantage of a bug in such a way that just viewing an e-mail message caused the virus to infect your computer! And we’re *paying* for this quality of software?

Note Although I’m being disparaging about Windows here, unlike many books, *Beginning Ubuntu Linux* doesn’t ignore Windows. Throughout its pages, you’ll find frequent references to Windows and the software that runs under it. You’ll find direct comparisons with actual Windows programs. The intention is that anyone with prior experience will be able to get started with Ubuntu much more quickly.

And how about the speed at which Windows runs? It's just dandy when your PC is brand new. But after just a few months, it seems like someone has opened up the case and poured molasses inside. It takes quite a few seconds for My Computer to open, and there's time for a coffee break while Internet Explorer starts up.

So is Linux the solution to these problems? Most would agree that it's a step in the right direction, at the very least. Linux doesn't need antivirus programs, because there are virtually no Linux-specific viruses. As with all software, security holes are occasionally discovered in Linux, but the way it is built means exploiting those holes is much more difficult.

Note There have been a couple of viruses for Linux, but they're no longer "in the wild" (that is, they are no longer infecting PCs). This is because the security holes they exploited were quickly patched, causing the viruses to die out. This happened because the majority of Linux users update their systems regularly, so any security holes that viruses might exploit are patched promptly. Compare that to Windows, where most users aren't even aware they can update their systems, even when Microsoft gets around to issuing a patch (which has been known to take months).

There's also the fact that Linux encourages you to take control of your computer, as opposed to treating it like a magical box. As soon as you install Linux, you become a power user. Every aspect of your PC is under your control, unlike with Windows. This means fixing problems is a lot easier, and optimizing your system becomes part and parcel of the user experience. You no longer have to take poor performance lying down. You can do something about it!

WINDOWS COMPATIBLE?

One of the biggest questions asked by most newcomers to Linux is whether it can run Windows software. The answer is yes . . . and no.

Linux is completely different from Windows on a fundamental technical level. Its creators based it on Unix, an industrial-strength operating system, and deliberately steered clear of emulating Windows. This means that Linux isn't a swap-in replacement for Windows. You cannot take the installation CD of a Windows program and use it to install that program on Linux, in the same way that you cannot install an Apple Mac program on Windows.

However, several current projects let you run Windows programs on Linux. Wine (www.winehq.com) is an example of such a project, and you can download a commercial and easy-to-use variation of it from www.codeweavers.com. You can also use programs like VMware (www.vmware.com) to create a "virtual PC" running on Linux. Then you can install the Windows operating system and, therefore, any Windows software you like.

In most cases, however, you'll find that there's a Linux equivalent of your favorite Windows software. Frequently, you'll find that this Linux version is actually superior to the Windows program you've been using. I'll discuss many of these in Chapter 11.

The Benefits of Linux

People have been known to exaggerate about Linux when singing its praises, and there's certainly some hyperbole around. But there are a couple of cast-iron facts about its benefits.

Crash-Free

A primary benefit of Linux is that it doesn't crash. In years and years of using Linux, you will never experience your mouse cursor freezing on screen. A strange error box won't appear and not go away until you reboot. It's possible to leave a Linux system running for years without ever needing to reboot (although most desktop Ubuntu users shut down their PC when they won't be using it for a while, just like the rest of us).

Of course, programs that run on top of Linux sometimes crash, but they don't take the rest of the system down with them, as can happen under Windows. Instead, you can clean up after a crash and just carry on.

Note Actually, very few programs under Linux crash. Because Linux programmers use a different method of bug testing than used by Microsoft developers, there are arguably fewer bugs, and those that are discovered are fixed very quickly.

Security

The next benefit is that Linux is far, far more secure than Windows. Linux is based on years of proven computer science research. It works on the principle of users who have permissions to undertake various tasks on the system. If you don't have the correct permission, then you cannot, for example, access a particular piece of hardware. Additionally, privacy can be ensured because the files on the PC are "owned" by individual users, who can permit or deny others access to those files.

Free and Shareable

Another big benefit is that Linux can be obtained free of charge. Once it's installed, the latest updates for all your programs are also free of charge. Not only that, but if you want any new software, it will also usually be free of charge (and normally just a download away). Is this starting to sound attractive yet?

Because the software is free, you can share it with friends. Suppose that you find a really great image editor. You mention it to a friend, and he asks for a copy. Under Windows, copying the program is strictly illegal—to do so turns you into a software pirate! Unless that image editor is freeware, your friend will need to buy the software himself. Under Linux, sharing software is normally entirely legal. In fact, it's encouraged! I'll explain why in Chapter 2.

Compatible with Older Hardware

Another benefit of Linux is that it works very well on older hardware and doesn't require a cutting-edge PC system. The latest version of Windows XP requires high-powered hardware,

to the extent that upgrading to that operating system usually means buying a new PC, even if your old one still works fine!

In contrast, Linux works on computers dating back as far as the early 1990s. This book was largely written on a five-year-old Pentium II 450 MHz notebook running Ubuntu. Although it would be an exaggeration to claim that the computer is lightning-fast, there's little waiting around for programs to start. On the same machine, Windows 2000 (which came installed on the computer) grinds and churns, and using it can be a frustrating experience.

Linux encourages an attitude of both recycling and making the most of what you have, rather than constantly upgrading and buying new hardware. You can pull out that "old" PC and bring it back to life by installing Linux. You might even be able to give it away to a family member or friend who does not have a PC. Perhaps it's time for grandma to get online, or perhaps you can give the kids their own PC so they will stop using yours.

Alternatively, you might consider turning old hardware into a server. Linux is capable of just about any task. As well as running desktop computers, it also runs around 60% of the computers that make the Internet work. Linux is extremely flexible. You could turn an old PC into a web server, e-mail server, or firewall that you can attach to a broadband Internet connection. If you were to do this with Microsoft software, it would cost hundreds of dollars, not to mention requiring an advanced computer. It's free with Linux.

The Linux Community

So we've established that Linux is powerful, secure, and flexible. But I've saved the best for last. Linux is more than a computer operating system. It's an entire community of users spread across the globe. When you start to use Linux, you become part of this community (whether you like it or not!).

One of the benefits of membership is that you're never far from finding a solution to a problem. The community likes to congregate online around forums and newsgroups, which you can join in order to find help.

Your placement in the ranks of the community is "newbie." This is a popular way of describing someone who is new to Linux. Although this sounds derisory, it will actually help when you talk to others. Advertising your newbie status will encourage people to take the time to help you. After all, they were newbies once upon a time!

There's another reason not to be disheartened by your newbie tag: you'll outgrow it very quickly. By the time you reach the end of this book, you'll have advanced to the other end of the spectrum—"guru." You'll be one of those giving out the advice to those poor, clueless newbies, and you'll be 100% confident in your skills.

Tip One of the best ways to learn about Linux is under the auspices of a knowledgeable friend. It's very beneficial to have your own guru to help you along when you get stuck—someone who is just an e-mail message or phone call away. If you have a friend who uses Linux, consider taking him or her out for a drink and getting more friendly!

But being part of a community is not just about getting free technical support. It's about sharing knowledge. Linux is as much about a political ideal as it is about software. It was created to be shared among those who want to use it. There are no restrictions, apart from one: any changes you make must also be made available to others.

The spirit of sharing and collaboration has been there since day one. One of the first things Linus Torvalds did when he produced an early version of Linux was to ask for help from others. And he got it. Complete strangers e-mailed him and said they would contribute their time, skills, and effort to help his project. This has been the way Linux has been developed ever since. Hundreds of people around the world contribute their own small pieces, rather than there being one overall company in charge. And the same concept applies to knowledge of Linux. When you learn something, don't be afraid to share this knowledge with others. "Giving something back" is a very important part of the way of Linux.

To understand why Linux is shared, you need to understand its history, as well as the history of what came before it. This is the topic of Chapter 2.

Summary

This chapter provided an introduction to Linux. It explained what Linux can be used for and also its many advantages when compared to Microsoft Windows. It also introduced the community surrounding Linux, which adds to its benefits. You should be starting to realize what makes millions of people around the world use Linux as the operating system of choice.

The next chapter covers the history of Linux. It also discusses another curious aspect: the political scene that drives the operating system forward.



A History and Politics Lesson

Linux is more than just software. It's an entire community of users, and as such, there's a detailed social history behind it. In this chapter, we'll look at the origins of Linux, both in terms of where it came from and the people who make it.

You might be tempted to skip this chapter and move on to the information about installing Ubuntu. To be fair, nothing of vital technical importance is mentioned here. But it's important that you read this chapter at some stage, because Linux is more than simply the sum of its parts. It's far more than simply a set of computer programs.

If nothing else, this chapter explains the fundamental philosophies behind Linux and attempts to answer some of the often-baffling questions that arise when Linux is considered as a whole.

In the Beginning

Linux was created 15 years ago, in 1991. A period of 15 years is considered a lifetime in the world of computing, but Linux actually harks back even further, into the early days of modern computing in the mid-1970s.

Linux was created by a Finnish chap named Linus Torvalds. At the time, he was studying in Helsinki and had bought a desktop PC. His new computer needed an operating system. Torvalds's operating system choices were limited: there were various versions of DOS and something called Minix. It was the latter that Torvalds decided to use.

Minix was a clone of the popular Unix operating system. Unix was used on huge computers in businesses and universities, including those at Torvalds's university. Unix was created in the early 1970s and has evolved since then to become what many considered the cutting edge of computing. Unix brought to fruition a large number of computing concepts in use today and, many agree, got almost everything just right in terms of features and usability.

Versions of Unix were available for smaller computers like Torvalds's PC, but they were considered professional tools and were very expensive. This was in the early days of the home computer craze, and the only people who used IBM PCs were businesspeople and hobbyists.

Note Linux is a pretty faithful clone of Unix. If you were to travel back in time 20 or 30 years, you would find that using Unix on those old mainframe computers, complete with their teletype interfaces, would be similar to using Linux on your home PC. Many of the fundamental concepts of Linux, such as the file system hierarchy and user permissions, are taken directly from Unix.

Torvalds liked Unix because of its power, and he liked Minix because it ran on his computer. Minix was created by Andrew Tanenbaum, a professor of computing, to demonstrate the principles of operating system design to his students. Because Minix was also a learning tool, people could also view the *source code* of the program—the original listings that Tanenbaum had entered to create the software.

Minix was lacking in some significant areas. Many people, including Torvalds, found using it very frustrating. Torvalds decided to create from scratch his own version of Minix, but to make it better, avoiding what many considered the pitfalls of Minix. He managed to produce version 0.01 of Linux in just over half a year.

Note Most clones or implementations of Unix are named so that they end in an *x*. One story has it that Torvalds wanted to call his creation Freax, but a containing directory was accidentally renamed Linux on an Internet server. The name stuck.

From day one, Torvalds intended his creation to be shared among everyone who wanted to use it. He encouraged people to copy it and give it to friends. He didn't charge any money for it, and he also made the source code freely available. The idea was that people could take the code and improve it.

This was a master stroke. Many people contacted Torvalds, offering to help out. Because they could see the program code, they realized he was onto a good thing. Soon, Torvalds wasn't the only person developing Linux. He became the leader of a team that used the fledgling Internet to communicate and share improvements.

Note The popular conception of Linux is that it is created by a few hobbyists who work on it in their spare time. This might have been true in the very early days. Nowadays, in addition to these “bedroom programmers,” Linux is programmed by hundreds of professionals around the world, many of whom are employed specifically for the task. Torvalds adds to the effort himself and also coordinates the work.

It's important to note that when we talk here about Linux, we're actually talking about the kernel—the central program that runs the PC hardware and keeps the computer ticking. This is all that Torvalds initially produced back in 1991. It was an impressive achievement, but needed a lot of extra add-on programs to take care of even the most basic tasks. Torvalds's kernel needed additional software so that users could enter data, for example. It needed a way for users to be able to enter commands so they could manipulate files, such as deleting or copying them. And that's before you even consider more complicated stuff like displaying graphics on the screen or printing documents.

Linux itself didn't offer these functions. It simply ran the computer's hardware. Once it booted up, it expected to find other programs. If they weren't present, then all you saw was a blank screen.

LINUS TORVALDS

Linus Benedict Torvalds was born in Helsinki, Finland, in 1969. A member of the minority Swedish-speaking population, he attended the University of Helsinki from 1988 to 1996, graduating with a Masters degree in Computer Science.

He started Linux not through a desire to give the world a first-class operating system but with other goals in mind. Its inspiration is in part due to Helsinki winters being so cold. Rather than leave his warm flat and trudge through the snow to the university's campus in order to use its powerful minicomputer, he wanted to be able to connect to it from home! He also wanted to have a platform to use to experiment with the properties of the Intel 386, but that's another story. Torvalds needed an operating system capable of such tasks. Linux was born.

It took Torvalds the better part of a year to come up with the very first version of Linux, during which he worked alone in a darkened room. In 1991, he announced his creation to the world, describing Linux as “just a hobby,” and saying it would never be big. It wouldn't be until 1994 that it reached version 1.0.

In the early days, Torvalds's creation was fairly primitive. He was passionate that it should be free for everyone to use, and so he released it under a software license that said that no one could ever sell it. However, he quickly changed his mind, adopting the GNU Public License.

Torvalds was made wealthy by his creation, courtesy of the dot.com boom of the late 1990s, even though this was never his intention; he was driven by altruism. Nowadays, he lives in Portland, Oregon, with his wife and children, having moved to the United States from Finland in the late 1990s.

Initially, Torvalds worked for Transmeta, developing CPU architectures as well as overseeing kernel development, although this wasn't part of his official work. He still programs the kernel, but currently he oversees the Open Source Development Lab, an organization created to encourage open source adoption in industry and which is also referred to as the home of Linux.

The GNU Project

Around the time Linus created Linux, another project, called GNU, also existed. This project team also hoped to create an operating system that used Unix as its inspiration, although avoiding some of the pitfalls that had blighted that operating system, both technically and in terms of its licensing. GNU is a so-called recursive acronym that stands for “GNU's Not Unix,” a play on words favored by computer programmers.

GNU's parent organization, the Free Software Foundation (FSF), had been formed eight years prior to Torvalds's effort, and since that time, had produced the majority of the core software that Linux desperately needed. However, as luck would have it, FSF lacked the essential functionality of the kernel. The developers were in the process of creating their own kernel, but it had not come to fruition.

The GNU software was distributed for free to anyone who wanted it. The source code was also made available so users could adapt and change the programs to meet their own needs (in fact, Torvalds had used the GNU model when deciding how to distribute Linux).

Richard Stallman is the man behind GNU and, along with Linus Torvalds, is the second accidental hero in our story. Stallman had been around since the Dark Ages of computing, back when wardrobe-sized computers were “time-shared” among users who used small desktop terminals to access them. Like Torvalds, Stallman started GNU as a personal project, but then found others who were more than willing to join his cause.

Note Stallman created the Emacs text editor and the GNU C Compiler (GCC). Together, they allow the creation of yet more software, so it's no surprise that one of the very first programs Torvalds used in the early days to create Linux was Stallman's GCC.

Back in Stallman's day at the legendary Massachusetts Institute of Technology (MIT), computer software was shared. If you came up with a program to perform a particular task, you offered it to practically anyone who wanted it. Alternatively, if you found an existing program wasn't adequate or had a bug, you improved it yourself, and then made the resulting program available to others. People might use your improved version, or they might not; it was up to them.

This way of sharing software was disorganized and done on an ad hoc basis, but came about of its own accord. Nobody questioned it, and it seemed the best way of doing things. There certainly wasn't any money involved, any more than there would be money involved in one friend explaining an idea to another.

RICHARD STALLMAN

Richard Matthew Stallman, usually referred to as RMS, was born in 1953 in Manhattan. He comes from the old school of computing forged during the 1970s and was a member of MIT's legendary Artificial Intelligence Lab.

Seemingly destined for a life in academia, Stallman left MIT in 1984 to found the GNU Project. This was as a reaction to the increasing commercialization of computer software. Whereas once all hackers (that is, programmers) had shared ideas and program code, the trend in the 1980s was toward proprietary, nonshared code, as well as legal contracts, which forced programmers to keep secrets from one another.

Stallman is a very talented programmer and is considered a genius by many observers. He single-handedly created many essential programming tools in his initial efforts to get GNU off the ground. Many of these find a home in Linux.

Stallman is also widely applauded for the creation of the GNU Public License. This is a legal document that lets people share software. It introduces the concept of *copyleft* and is opposed to the legal concept of copyright, which attempts to limit the freedom of individuals when using a piece of software (or any other creative work). Nowadays, the concept of copyleft has been applied to literature, music, and other arts in an attempt to avoid restricting who can and cannot access various items, as well as to encourage a collaborative working environment.

Proprietary Software and the GPL

In the 1980s, everything changed. The world became more corporate, and with the rise of the desktop PC, the concept of proprietary software became prevalent. More and more companies started to sell software. They reasoned that this was impossible to do if they shared it with everybody else, so they kept it secret. Microsoft led this charge and did very well with its proprietary software.

To Stallman, this "trade secrets" approach to software was anathema. He had nothing against software being sold for a profit, but he hated the fundamental ideas behind software being kept secret. He felt passionately that sharing software and being able to understand how

it worked was akin to free speech—necessary and vital for the furthering of technology, and therefore society itself. How could the new generation of programmers improve on the previous generation’s work if they were unable to see how it worked? It was absurd to need to create software from scratch each time, rather than taking something that already existed and making it better.

Because of his beliefs, Stallman resigned from his job in the MIT Artificial Intelligence Lab and founded GNU. His aim initially was to produce a complete clone of Unix that would be shared in the ways he knew from the early days of computing. This software would be available for everyone to use, to study, and to adapt. It would be free, in the same sense as free speech—shared and unrestricted. This gave rise to the vital concept of “free software” and soon GNU, and the FSF, became not just a programming venture, but also a political movement.

Note A very common misconception of “free software” is that it is always free of charge. This isn’t correct. The word *free* is used here in its political sense, as in “free speech.” Many companies and individuals make a healthy profit from selling free software and, in fact, selling free software is encouraged by the GNU Project.

To protect the rights of people to share and adapt the GNU software, Stallman came up with the GNU Public License (GPL). Various drafts of this license were produced over time, until it became a completely watertight legal contract, which furthered the concept of free software.

Most software you buy comes with a license agreement—that big chunk of text you must agree to when installing software (in the case of Windows desktop software, it’s frequently referred to as the End-User License Agreement, or EULA). The license agreement usually says that you cannot copy the software or share it with friends. If others want to use the software, they must buy their own version.

The GPL turns this on its head. Rather than restricting what people can do with the software, it gives them permission to share the software with whomever they wish. However, if they modify the program in any way, and then distribute it to others, the program they come up with must also be licensed under the GPL. In other words, people cannot make changes to a program that has a GPL, and then sell the modified program, keeping their improvements secret.

Note An interesting side note is that the actual wording of the GPL says that any changes you make should be shared with others *only if the software is redistributed*. This means that if you modify some GPL software and don’t give it to anyone else, there’s no need for you to publish your changes or make others aware of those changes.

GNU and Linux Together

The Linux kernel, developed by Torvalds, and the GNU software, developed by Stallman, were a perfect match. It’s important to note that this doesn’t mean the two projects joined forces. It simply means that the Linux project took some of the GNU software and gave it a good home.

This was done with Stallman's blessing, but there wasn't any official union between the two groups. Remember that Stallman had intended everyone to freely share and use the GNU tools. Linux represented a set of people doing just that. GNU is still working on its own kernel, called Hurd, which may provide an alternative to using Torvalds's Linux kernel.

Note Hurd was first planned back in the 1980s and, at the time of writing, still has yet to see the light of day (although testing versions are available). Hurd is a hugely ambitious project and will set a gold standard when it is released.

GNU and Linux together formed a complete operating system, which mimicked the way Unix operated. Other projects and individuals spotted the success of Linux and came onboard, and it wasn't long before Linux realized the potential for a graphical user interface (GUI), the fundamentals of which were provided by the XFree86 Project. A lot of additional software was also provided by individuals and organizations, all using the same "share and share alike" example set by Stallman, with the GNU tools, and Torvalds, with his kernel.

Many people refer to Linux as GNU/Linux. This gives credit to the GNU Project that provided the majority of tools vital to making Linux into a usable operating system. However, like the majority of people in the computing world, I use the term *Linux* throughout this book to avoid confusion.

Different Flavors of Linux

All the pieces of GNU software were available for free download and were therefore free of charge. But this brought its own problems. Not everyone had the know-how to put all the bits and pieces together into a complete operating system. Those who could do this didn't necessarily have the time for it.

Because of this, a number of companies stepped in to do the hard work. They put together versions of Linux, complete with all the software from the GNU Project, which they then sold for a fee on floppy disks, CDs, or DVDs. They also added in bits of their own software, which made it possible to install Linux easily onto a computer's hard disk, for example. They produced their own manuals and documentation, too, and did other things such as bug testing to ensure it all worked well.

What they came up with became known as *distributions* of Linux, or *distros* for short. Examples of these companies include Red Hat, SUSE, Mandrake, and many others around the world. Additionally, a number of enthusiasts got together and formed organizations to create their own distros, such as Debian and Slackware.

Modern distros are very advanced. They make it easy to install Linux on your PC, and they usually come with everything you need, so you can get started immediately. Additionally, they have their own look and feel, as well as unique ways of working and operating. This means that Ubuntu is not the same as Red Hat Linux, for example, although they share a lot of common features and, of course, they all share the core GNU software.

Linux Today

Nowadays, Linux is a thoroughly modern and capable operating system, considered cutting edge by many. It also runs on many different types of computer hardware, including Apple Macintosh computers, Sun SPARC machines, and the ubiquitous desktop PCs equipped with Intel or AMD processors. One of the ironies is that, although Linux was based on Unix, it has slowly come to dominate the computer operating system market. According to industry sources, Linux is on its way to making commercial varieties of Unix redundant. Companies that sell their own versions of Unix, such as Hewlett Packard and IBM, have added Linux to their traditional product range.

Recent innovations in the latest versions of the kernel mean that it finds uses on the smallest computers in the world, as well as on the biggest. Several of the top supercomputers in the world run Linux and, ironically, it can also be used on handheld PDAs or even digital watches! You'll even find it running things like digital video recorders or other household goods, where it sits invisibly in the background and makes everything work. Remember that one of the fundamental principles of Linux is that you can use it for whatever you want. You don't need to ask for permission first or tell anyone what you're doing.

Linux initially found mainstream use by software developers, and on server computers, such as those that run the Internet. However, in recent years, it has become increasingly popular on desktop computers. This is the area where experts suggest it will see massive growth over the coming years.

Modern Linux Development

Nowadays, Linux is developed not only by Torvalds, who manages the huge project, but also by hundreds of volunteers and corporations who contribute resources. Most recently, IBM and Novell have gotten involved and contribute hundreds of people to the effort of creating Linux. Sun contributes the OpenOffice.org office suite and sells its own version of Linux. Corporations like Computer Associates contribute their own software, too.

These companies have realized that the best way of producing software is to share and share alike, rather than develop their own proprietary software and keep it secret. The proprietary ways of the 1980s are starting to seem like an ill-conceived flash in the pan.

Most recently, Novell found that by embracing Linux, it could massively enhance the functions of its aging NetWare product, yet without needing to return to the drawing board and start from scratch. It could just take what it wanted from the pile of Linux software. This shows the philosophy of Linux in action.

Linux has software for just about every need, ranging from simply receiving e-mail to running a huge e-mail server. There are databases, office suites, web browsers, video games, movie players, audio tools, and more, as well as thousands of pieces of specialized software used in various niches of industry (and too boring to mention here). Most of this software is available to anyone who wants it, free of charge.

What more could you want?

Summary

This chapter has detailed the history of Linux and explained its origins. It also explained *why* Linux came into being. We looked at how Linux formed one of the building blocks of a political movement geared toward producing software that can be shared.

We discussed the creator of Linux, Linus Torvalds. We've also looked at the massive input the GNU Project has made and, in particular, that of its philosopher king, Richard Stallman.

In the next chapter, we move on to look at what you can expect from day-to-day use of Linux.



The Realities of Running Linux

So now that you've learned about the politics, history, and personalities behind Linux, only one question remains: what's Linux actually like when used day to day? What should the average user expect from the experience?

These are the questions I hope to answer in this brief chapter.

Learning to Use Linux

What should you expect from Linux once you've installed it? Well, it's a little like running Windows, except there are no viruses, fewer crashes, and no inexplicable slowdowns.

In addition, you have complete control over the system. This doesn't mean Linux is necessarily complicated. It's just that you have the control if you wish to make use of it. We'll look into this in the later chapters of this book.

Most software you use under Windows has at least one equivalent under Ubuntu, installed by default. It's unlikely that you'll need to download or install any additional software and, even if you do, you'll probably find it's available for free.

In most cases, the Linux swap-ins are at least as powerful and easy to use as their Windows alternatives. Tabbed browsing in the Mozilla Firefox web browser lets you visit more than one site at once, for example, without needing to have a lot of browser instances running, as you do with Microsoft Internet Explorer. The Evolution program has a search routine that lets you look through your e-mail messages quickly for a variety of criteria, and it puts the features in a similar Microsoft product to shame.

Does this sound too good to be true? There is just one caveat. Linux isn't a clone of Windows and doesn't aim to be. It has its own way of doing certain things, and sometimes works differently from Windows. This means that many people experience a learning curve when they first begin using Linux.

Note Several Linux distributions aim to mimic Windows pretty faithfully. For example, Xandros, Lycoris/lx, and Linspire copy the look and feel of Windows to the extent that (allegedly) some people are unable to tell the difference.

But in just a few weeks after your move to Linux, everything will start to seem entirely normal. Most of the time, you won't even be aware you're running Linux. Of course, some patience is required during those initial few weeks. Linux can be illogical and frustrating; on the other hand, so can Windows. We simply got used to it.

Who Uses Linux?

Who uses Linux? The myth from the old days is that it's only for techies and power users. When you needed to put everything together by hand, this was clearly true. But modern distributions make Linux accessible to all. It's no exaggeration to say that you could install Linux on a computer Luddite's PC and have that person use it in preference to Windows.

Up until quite recently, Linux was largely seen as a developer's tool and a server operating system. It was geared toward programmers or was destined for a life running backroom computers, serving data, and making other computer resources available to users.

To this end, Linux continues to run a sizable proportion of the computers that make the Internet work, largely because it provides an ideal platform for the Apache web server, as well as various databases and web-based programming languages. This has led to the LAMP acronym, which stands for Linux, Apache (a web server), MySQL (a database), and PHP, Python, or Perl (three programming languages that can be used in an online environment).

Despite its technical origins, recent years have seen a strong push for Linux on desktop computers. Linux has stepped out of the dark backrooms, with the goal of pushing aside Microsoft Windows and Mac OS in order to dominate the corporate workstation and home user market.

Running Linux on the desktop has always been possible, but the level of knowledge required was often prohibitively high, putting Linux out of the reach of most ordinary users. It's only comparatively recently that the companies behind the distributions of Linux have taken a long, hard look at Windows and attempted to mirror its user-friendly approach. In addition, the configuration software in distributions like Ubuntu has progressed in leaps and bounds. Now, it's no longer necessary to know arcane commands in order to do something as simple as switch the screen resolution. The situation has also been helped by the development of extremely powerful office software, such as OpenOffice.org and Koffice.

Is Linux for you? There's only one way of finding out, and that's to give it a go. Linux doesn't require much of you except an open mind and the will to learn new ways of doing things. You shouldn't see learning to use Linux as a chore. Instead, you should see it as an adventure—a way of finally getting the most from your PC and not having to worry about things going wrong for reasons outside your control.

Linux puts you in charge. You're the mechanic of the car as well as its driver, and you'll be expected to get your hands dirty every now and then. Unlike with Windows, Linux doesn't hide any of its settings or stop you from doing things for your own protection; everything is available to tweak. Using Linux requires commitment and the realization that there are probably going to be problems, and they're going to need to be overcome.

However, using Linux should be enjoyable. In his initial newsgroup posting announcing Linux back in 1992, Linus Torvalds said that he was creating Linux "just for fun." This is what it should be for you.

Getting Hold of Linux

Getting hold of Linux is easy. You'll already have spotted the version of Ubuntu packaged with this book. Ubuntu is the main focus of this book, and I consider it to be one of the very best versions of Linux out there. It's ideal for both beginners and power users, and it really does match the functionality offered in Windows. It includes several easy to use configuration tools, which makes changing your system settings a breeze. For example, the Synaptic Package Manager software can automate the download and installation of new software with just a few clicks.

Ubuntu is also a very good-looking distribution. You'll find your friends and colleagues “wowing” when they happen to pass by and glance at your PC!

Quite a number of Linux distributions are available. If you want to explore other Linux distributions as well as Ubuntu, by far the most fuss-free method of getting hold of Linux is to pop over to your local computer store (or online retailer) and buy a boxed copy. You can choose from Red Hat, SUSE, Mandrake, Libranet, TurboLinux (if you want foreign language support, although nearly all commercial distributions do a good job of supporting mainstream languages), and many others. Many distributions come on more than a single CD—typically up to four CDs at the moment. Some versions of Linux come on DVD.

Caution Bearing in mind what I've said about the sharing nature of Linux, you might think it possible to buy a boxed copy of Linux and run off copies for friends, or even sell them for a profit. However, you shouldn't assume this is the case. A minority of distribution companies, such as Xandros and Linspire, incorporate copyrighted corporate logos into their distributions that place restrictions on redistribution. Sometimes they include proprietary software along with the Linux tools, which you cannot copy without prior permission. However, in many cases, reproducing the CDs in small volumes for friends or for use on workstations in a company environment is permitted.

Many of the Linux distributions are also available to download free of charge. In fact, many community-run distributions—such as Slackware, Debian, Fedora, and Gentoo—are *only* available this way (although you can often buy “homemade” CDs from smaller retailers, who effectively burn the CDs for you and produce makeshift packaging). If your PC has a CD-R/RW drive and you have some CD-burning software under Windows (such as Nero), you can download an ISO image and make your own installation CD from it.

Note An *ISO image* is a very large file (typically 700MB), which you can burn to CD. This CD is then used to install Linux.

Using Ubuntu

Ubuntu is a relatively young Linux distribution. It is based on and is still closely allied with the Debian distribution, like many versions of Linux. Debian has been in existence for almost as long as the Linux kernel and embraces the spirit and philosophy of Linux, which says software should be shared and made available to anyone who wants it.

Ubuntu takes this one step further. Its goal is to give anyone in the world access to an easy-to-use version of Linux, regardless of geographical location or physical abilities. Ubuntu supports a large number of languages, so it can be used in most countries around the world. In addition, it includes accessibility tools, so it can be used by partially sighted, deaf, or disabled people.

From the very start, Ubuntu's creators decided it would always be free of charge and would always be freely available. Unlike with many versions of Linux, no commercial version of Ubuntu exists (although it is possible to pay for various services, such as support, if you wish).

But perhaps Ubuntu's greatest strength is its community, which extends across the world. If you have a question about Ubuntu, you'll find hundreds of people willing to help. Just as the software is designed to be shared, a strong belief within the Ubuntu community is that knowledge should be shared, too.

It might come as no surprise that *Ubuntu* is an ancient African word that roughly translates as "humanity to others." In South Africa in particular, "Ubuntu" is a way of life that advocates acceptance and compassion towards others. The Ubuntu philosophy is one of the underpinnings of post-apartheid South Africa and spreads into political spheres as well as everyday life. The Ubuntu version of Linux is an attempt to bring this spirit to the world of computer software.

Ubuntu is primarily geared toward desktop users, although with a little adapting, it can also be used to run server computers.

Ubuntu is designed to be easy to use. Anyone who has used Windows or Mac OS will feel right at home. It features every piece of software you could wish for or would find within a well-equipped modern operating system. It includes a web browser, an e-mail client, instant messaging software, an office suite, a graphics editor, and much more. And don't think that these are cut-down versions designed to lure you into purchasing the full version later on. In every case, they're full-featured pieces of software that give proprietary programs a run for their money.

Perhaps more important, Ubuntu is very user-friendly. Updating the system can be done with just a few clicks of the mouse, as can downloading and installing new software.

Summary

This chapter explained what you can realistically expect when using Linux every day. It also discussed the kind of company you'll be keeping in terms of fellow users.

You learned how people usually get hold of Linux. Of course, with this book, you already have a version of Linux, Ubuntu, which was introduced in this chapter.

This completes the general overview of the world of Linux. In the next part of the book, you'll move on to actually installing Linux on your hard disk. This sounds more daunting than it actually is. The next chapter gets you started by explaining a few basic preinstallation steps.

PART 2



Installing Ubuntu



Preinstallation Steps

The first part of this book discussed the pros and cons of using Linux as part of your day-to-day life. It was intended to help you evaluate Linux and understand what you're buying into should you decide to make it your operating system of choice. Now, we move on to actually installing Linux and, specifically, Ubuntu, which is included with this book on a CD-ROM.

Installing any kind of operating system is a big move and can come as something of a shock to your PC. However, Ubuntu makes this complicated maneuver as easy as it's possible to be. Its installation routines are very advanced compared to previous versions of Linux, and even compared to other current distributions.

What does saying that you're going to install Ubuntu actually mean? This effectively implies three things:

- Somehow all the files necessary to run Ubuntu are going to be put onto your hard disk.
- The PC will be configured so that it knows where to find these files when it first boots up.
- The Ubuntu operating system will be set up so that you can use it.

However, in order to do all this and get Ubuntu onto your PC, you must undertake some preparatory work, which is the focus of this chapter.

Understanding Partitioning

Chances are, if you're reading this book, your PC already has Windows installed on it. This won't present a problem. In most cases, Ubuntu can live happily alongside Windows in what's called a *dual-boot setup*, where you can choose at startup which operating system to run. However, installing Ubuntu means that Windows must make certain compromises. It needs to cohabit your hard disk with another operating system—something it isn't designed to do.

The main issue with such a situation is that Windows needs to shrink and make some space available for Ubuntu (unless you install a second hard disk, which is discussed later in this chapter). Ubuntu isn't able to use the same file system as Windows, and it needs its own separately defined part of the disk, which is referred to as a *partition*. All of this can be handled automatically by the Ubuntu installation routine, but it's important that you know what happens.

All hard disks are split into partitions, which are large chunks of the disk created to hold operating systems (just like a large farm is partitioned into separate fields). A partition is usually multiple gigabytes in size, although it can be smaller. You can view your disk's partitions using the Disk Management tool in Windows XP and Windows 2000, as shown in Figure 4-1. You can

access this tool by opening Control Panel, clicking the Administrative Tools icon, selecting Computer Management, and then choosing Disk Management.

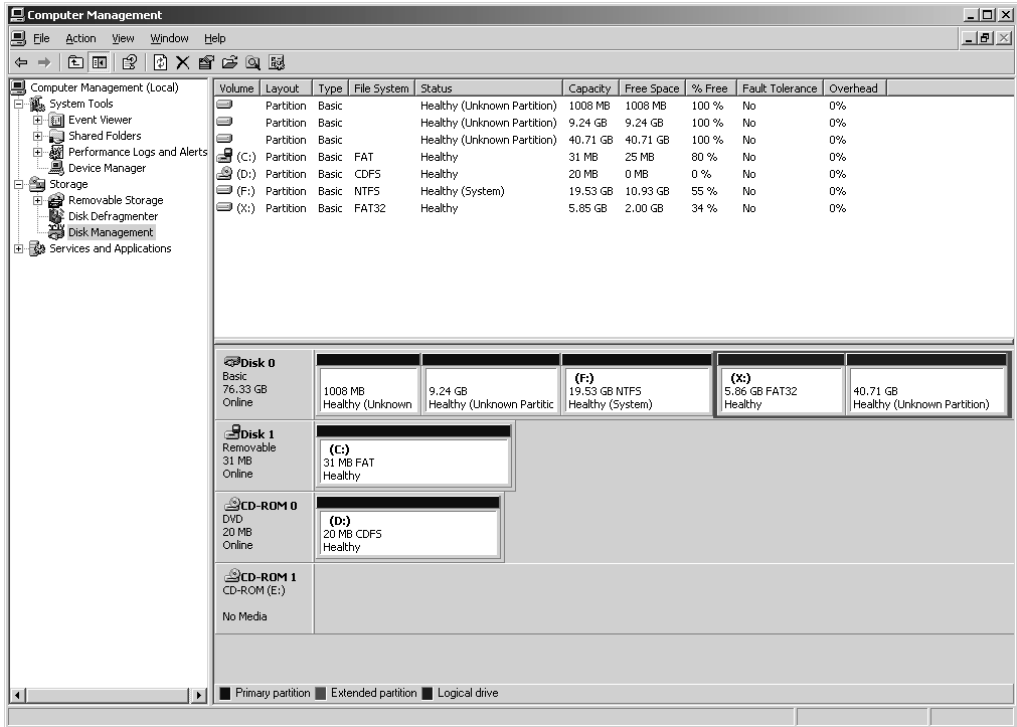


Figure 4-1. You can view your disk's partitions using Windows XP's Disk Management tool.

Most desktop PC systems have just one partition, unless the user has specifically created new partitions. As mentioned, Ubuntu needs a partition of its own. During installation, Ubuntu needs to shrink the main Windows partition and create a fresh partition alongside it (actually, it creates two partitions; the extra one is used to hold the swap file).

In addition, the Ubuntu installation routine writes a new *boot sector* (also known as a *boot loader*). The boot sector is located at the very beginning of the disk and contains a small program that then runs another program that lets you choose between operating systems (and therefore partitions) when you first boot up.

Note Not all Linux distributions have the ability to repartition the hard disk. In fact, at the time of writing, it's pretty rare. Most expect to simply take over the entire hard disk, wiping Windows in the process (although they'll always ask the user to confirm this beforehand). The ability to repartition a disk is just one of the reasons that Ubuntu is among the best Linux distributions currently available.

Of course, Ubuntu cannot shrink a Windows partition that is packed full of data, because no space is available for it to reclaim.

UBUNTU AND WINDOWS FILE SYSTEMS

One of the benefits of dual-booting Linux and Windows is that Ubuntu lets you access the files on the Windows partition. This is quite handy and facilitates the easy exchange of data.

If the Windows partition is FAT32—used on Windows 95, 98, Me, and (sometimes) XP—then Ubuntu can both read and write files to the partition. However, if the file system is NTFS—used with Windows NT, 2000, and (sometimes) XP—then Ubuntu will make the file system available as read-only. Because of this, you might consider converting your NTFS Windows partition to FAT32 before you install Ubuntu (but be aware that doing so means you lose some of the security and performance features of NTFS). Microsoft doesn't include a tool that lets you do this automatically, but you can use third-party disk partitioning programs like Symantec's Partition Magic (www.symantec.com) to convert your file system.

Freeing Up Space

The first step before installing Ubuntu alongside Windows is to check how much free space you have in your Windows partition. To see the amount of free space you have, double-click My Computer, right-click your boot drive, and select Properties. The free space is usually indicated in purple on a pie-chart diagram, as shown in Figure 4-2.




Figure 4-2. Ubuntu needs free disk space in which to install, so you might need to clean up your Windows partition.

You need to have at least 2GB of free space in your Windows partition for Ubuntu to use. You'll need more space if you wish to install a lot of programs. If you don't have enough free space, you have several options: reclaim space, remove Windows, or use a second hard disk.

Reclaiming Space

On Windows XP, you can run the Disk Cleanup tool to free some space on your hard disk. Click the Disk Cleanup button beneath the pie-chart diagram showing the free disk space (see Figure 4-2). Disk Cleanup is also accessible by choosing Start ► All Programs ► Accessories ► System Tools ► Disk Cleanup.

You might also consider turning off System Restore. This consumes a lot of disk space, which you can therefore reclaim. However, deactivating System Restore will mean that you lose the possibility of returning your system to a previous state should anything go wrong. To access the System Restore control, right-click My Computer, click Properties, and then click the System Restore tab.

If you still cannot free up enough disk space, consider uninstalling unused software via the Add/Remove Programs applet within Control Panel. If you have any large games installed, consider removing them first, because they usually take up substantial amounts of hard disk space. You might also consider deleting movie and MP3 music files, which are renowned for eating up hard disk space. The average MP3 is around 4MB, for example, and one minute of video typically takes up 1MB of disk space!

Removing Windows

Some users might prefer a second, more radical option: getting rid of Windows completely and letting Ubuntu take over the entire hard disk. If you feel confident that Ubuntu will fulfill your needs, this is undoubtedly the most straightforward solution. You'll be able to do this during installation. However, this will also mean that any personal data you have will be lost, so you should first back up your data (as described shortly).

Caution You should be aware that installing Windows back onto a hard disk that has Ubuntu on it is troublesome. Windows has a Darwinian desire to wipe out the competition. If you attempt to install Windows on a Ubuntu hard disk, it will overwrite Linux.

Using Another Hard Disk

A third option for making room for Ubuntu is attractive and somewhat safer in terms of avoiding the potential for data loss, but also potentially expensive: fitting a second hard disk to your PC. You can then install Ubuntu on this other hard disk, letting it take up the entire disk. Unlike Windows, Ubuntu doesn't need to be installed on the primary hard disk and is happy on a secondary drive.

A second hard disk is perhaps the best solution if you're low on disk space and want to retain Windows on your system. However, you'll need to know how to install the new drive or find someone to do it for you (although step-by-step guides can be found on the Web—just search

using Google or another search engine). In addition, if your PC is less than 12 months old, there is a possibility that you'll invalidate your guarantee by opening up your PC.

If you have an old PC lying around, you might also consider installing Ubuntu on it, at least until you're sure that you want to run it on your main PC. As noted in Chapter 1, one of the best features of Ubuntu is that it runs relatively well on older hardware. For example, a Pentium II with 256MB of memory should allow for very good performance.

NO-INSTALL LINUX

If you want to use the Linux operating system but leave your hard disk untouched, you might consider a number of additional options. Perhaps the most popular is to use a “live” version of Linux. A live version of Linux is one that boots and runs entirely from a CD (or DVD); it doesn't touch the user's hard disk. The most popular is Knoppix (www.knoppix.org). The distribution we discuss throughout this book, Ubuntu, also has a live CD version, which you can download from www.ubuntulinux.com (click the Download link and select the location closest to you).

Using Linux in this way involves downloading an ISO image (a single large file of around 650MB), which then must be burned to CD using a Windows program like Nero. Once the CD has been created, you simply boot from it in order to run Linux (after you ensure that your PC's BIOS is set to boot from CD, as explained in Chapter 5).

Alternatively, you might consider using virtual PC software. This type of software runs under Windows and re-creates an entire PC hardware system within software—effectively a PC within a PC. The hard disk is contained within one or two Windows files. Linux can then be installed on these virtual PC systems. When the program is switched to full-screen mode, it's impossible to tell you're running inside a computer system created in software. Two commercially available examples are considered worthwhile by many: VMware (www.vmware.com) and Microsoft's Virtual PC (www.microsoft.com/windowsxp/virtualpc/). You should be aware that both are designed to be professional-level tools, so they are quite expensive. You can also obtain open-source renditions of virtual PC software, such as QEMU (<http://fabrice.bellard.free.fr/qemu/>).

Another option in its infancy at the time of writing is Cooperative Linux, or coLinux for short (www.colinux.org). This is a set of Windows programs that aims to let Linux run under Windows using emulation. Unfortunately, setting up and using coLinux requires some expert knowledge, so you might want to wait until you have more experience with Linux.

Backing Up Your Data

Whichever route you decide to take when installing Ubuntu, you should back up the data currently on your computer beforehand. Possibly the easiest way of doing this is to burn the data to CD-R/RW discs using a program like Nero and a CD-R/RW drive.

If you take the coexistence route, installing Ubuntu alongside Windows, backing up your data should be done for insurance purposes. Although the Ubuntu Foundation tests all its software thoroughly and relies on community reporting of bugs, there's always the chance that something will go wrong. Repartitioning a hard disk is a major operation and carries with it the potential for data loss.

If you intend to erase the hard disk when installing Ubuntu (thereby removing Windows), you can back up your data, and then import it into Ubuntu.

Table 4-1 shows a list of common personal data file types, their file extensions, where they can be typically found on a Windows XP system, and notes on importing the data into Ubuntu. Note that earlier versions of Windows (95, 98, and Me) may differ when it comes to data storage locations.

Table 4-1. *Data That Should Be Backed Up*

Type of File	File Extensions	Typical Location	Notes
Office files	.doc, .xls, .ppt, .pdf, etc.	My Documents	Microsoft Office files can be opened, edited, and saved under Ubuntu using the OpenOffice.org suite. PDF documents can be viewed with the Evince program.
E-mail files	N/A	N/A	The Evolution mail client used by Ubuntu cannot import data directly from Microsoft Outlook or Outlook Express. However, there is a convoluted but effective workaround, which is described in the next section.
Digital images	.jpg, .bmp, .tif, .png, .gif, etc.	My Pictures (within My Documents)	Ubuntu includes a variety of programs to both view and edit image files.
Multimedia files	.mp3, .mpg, .avi	Various	With some additional downloads, discussed in Chapter 18, programs under Ubuntu can play MP3 music files and some movie file formats.
Internet Explorer Favorites	None	\Documents and Settings\ <username>\Favorites	Your Favorites list cannot be imported into Ubuntu, but the individual files can be opened in a text editor in order to view their URLs, which can then be opened in the Ubuntu web browser.
Miscellaneous Internet files	Various	Various	You might also want to back up web site archives or instant messenger chat logs, although hidden data such as cookies cannot be imported.

Backing Up E-mail Files

Microsoft e-mail cannot be easily imported into Ubuntu because Microsoft prefers to create its own proprietary file formats, rather than use open standards recognized by the rest of industry. Most e-mail programs use the MBOX format, and this is true of Ubuntu as well as programs created by the Mozilla Foundation (the organization behind the Firefox web browser). However, Microsoft uses its own DBX file format for Outlook Express and PST format for Outlook.

As a workaround, you can download and install the free Mozilla Thunderbird e-mail client (available from www.mozilla.com) on your Windows system. In Thunderbird, select **Tools** ► **Import** to import your messages from Outlook, Outlook Express, or even the popular Eudora mail client. You will then be able to back up Thunderbird's mail files and import them into Evolution under Ubuntu.

To find where the mail files are stored, in Thunderbird, select Tools ► Account Settings, and then look in the Local Directory box. Back up each file that corresponds to a folder within your mail program (for example, Inbox, Sent, and so on). Note that you only need to back up the files *without* file extensions. You can ignore the `.sbd` folders as well as the `.msf` files.

Making Notes

When you're backing up data, a pencil and paper come in handy, too. You should write down any important usernames and passwords, such as those for your e-mail account and other online services. You might want to write down the phone number of your dial-up connection, for example, or your DSL/cable modem technical settings. Figure 4-3 shows an example of some information you might want to record.

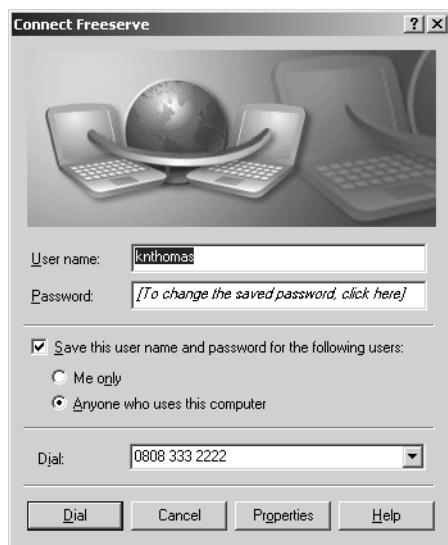


Figure 4-3. Don't forget to back up "hidden" data, such as Internet passwords.

In addition, don't forget to jot down essential technical details, such as your IP address if you are part of a network of computers using static addresses (this will usually be relevant only if you work in an office environment).

Tip If you've forgotten any passwords, several freeware/shareware applications are able to "decode" the asterisks that obscure Windows passwords and show what's beneath them. A good example is Asterisk Password Reveal, which you can download from www.paqtool.com/product/pass/pass_001.htm. Sites like www.download.com offer similar applications.

Note that you don't need to write down information such as hardware interrupt (IRQ) or memory addresses, because hardware is configured automatically by Ubuntu. However, it might

be worth making a note of the make and model of some items of internal hardware, such as your graphics card, modem (dial-up, DSL, or cable), and sound card. This will help if Ubuntu is unable to automatically detect your hardware, although such a situation is fairly unlikely to arise. You can garner this information by right-clicking My Computer on your desktop, selecting Properties, and then clicking the hardware tab. Then click the Device Manager button. Instead of writing everything down, you might consider taking a screenshot by pressing the Print Scr button and then using your favorite image editor to print it.

Tip Ubuntu works with a wide variety of hardware, and in most cases, it will automatically detect your system components. If you're in any doubt, you can consult the forums at www.ubuntuforums.org, in particular, the Hardware Help forums. You might also consider subscribing to one or more of the Ubuntu mailing lists at <http://lists.ubuntu.com>. Remember that an important element of Ubuntu is its community of users, many of whom will be very willing to answer any questions you might have!

Once you're certain that all your data is backed up, you can move on to the next chapter, which provides a step-by-step guide to installing the operating system.

Summary

The aim of this chapter has been to prepare both you and your computer for the installation of Ubuntu. We've looked at how your hard disk will be partitioned prior to installation and the preparations you should make to ensure your hard disk has sufficient free space. You also learned about the types of files you might choose to back up, in addition to vital details you should record, such as usernames and passwords for your online accounts.

In the next chapter, we move on to a full description of the Ubuntu installation procedure. The chapter guides you through getting Ubuntu onto your computer.



Installing Ubuntu

It's now time to install Linux. This chapter details how to install Ubuntu, the version of Linux supplied with this book. Ubuntu offers a special installation program for this purpose.

Installing Linux is a surprisingly quick task and shouldn't take more than 30 minutes on a modern PC. It's also relatively simple, with very few decisions to make throughout. Ubuntu's installation program automates the task to a high degree.

However, you should examine all the options you're offered to make sure they're correct. Installing an operating system involves a couple of serious maneuvers that, via an incorrect click of the mouse or accidental keystroke, bring with them the possibility of data loss. Read and consider every warning message you see, and be sure to keep your wits about you. Above all, make a backup of your data, as described in the previous chapter.

An Overview of the Installation Process

Installing Ubuntu requires little more of you than the ability to use the keyboard. The installation program is a little old-fashioned in that it doesn't use a mouse. Instead, you choose options by highlighting them using various keys and then pressing Enter. Don't worry—it's really easy, and you'll quickly get the hang of it. In fact, in most cases, the default choices are best, so throughout the entire procedure, you'll have to do little more than hit Enter a few times.

If you used MS-DOS in the old days, you'll have a head start when it comes to navigating around the screen. Pressing the Tab key moves the highlighted selection from option to option on the screen. Alternatively, you can use the arrow keys: press the right arrow to move the selection to the next option and the left arrow to move it to the previous one. Sometimes, you might need to move up and down in a list and, as you might expect, you can use the up and down arrow keys for that navigation!

When it comes to making a selection from a list, simply use the Tab or cursor keys to highlight the correct choice, and then press Enter. In the case of two or more choices, you might need to highlight each selection in turn, press the spacebar to choose it, and then move the highlight down to an option before hitting Enter. When you reach a list like this, how it works will be pretty obvious.

The installation works through a few stages in order to get Linux on your computer's hard disk. After a few preliminary setup options, you'll need to partition your hard disk in order to create space for Ubuntu. After this, Ubuntu is installed onto your hard disk. This is done in two phases. Initially, the software is copied onto your hard disk, and then, in a second automated phase, it's installed and set up on your computer.

Occasionally during the installation, you might have to answer a few questions. When the software that controls your monitor is installed, for example, you may be asked what resolution to use. You'll also need to set a username and password.

At the end of the procedure, your PC will boot straight into the Ubuntu login screen, and you're set to go. There's no need to mess around configuring hardware because that's done automatically. Neat, eh?

In most cases, the installation of Ubuntu will run smoothly without a hitch. But if you do run into problems, head over to Chapter 6, which addresses many of the most common issues and provides solutions.

Step-by-Step Guide

As outlined in Chapter 4, you shouldn't start the installation process until you've made sure there is enough space for Ubuntu on your hard disk and you have backed up all your data. With those preparations complete, you're ready to install Ubuntu. The remainder of this chapter guides you through the process.

Step 1: Boot from the CD-ROM

With your computer booted up, insert the Ubuntu disc into the CD/DVD-ROM drive. Close the tray, and then reboot your computer.

Ubuntu is installed by booting the installation program from CD-ROM. Therefore, the first step is to ensure your computer's BIOS is set to boot from the CD-ROM drive. How this is achieved depends on your particular computer, but it's always done via the BIOS setup program. Most PCs let you enter BIOS setup by pressing the Delete key just after the computer is first booted, although some use another key or key combination. Review the information displayed on your particular screen at boot time for more information.

When the BIOS menu appears, look for an option such as Boot and select it (you can usually navigate the BIOS menu by using the cursor keys and select options by pressing Enter). On the new menu, look for a separate entry such as Boot Device Priority or perhaps Boot Sequence. Ensure that the entry for the CD-ROM is at the top of the list, as shown in Figure 5-1. Arrange the list so that it's followed by the floppy drive and then your main hard disk (which will probably be identified as IDE-0 or First Hard Disk). You can usually press the F1 key for help on how the menu-selection system works.

Once you've made the changes, be sure to select the Exit Saving Changes option on the main BIOS setup menu. Your PC will then reset and boot from the Ubuntu CD-ROM, and you'll be greeted by the Ubuntu welcome screen, as shown in Figure 5-2. Press Enter to start the installation program.

Note In most cases, you can simply hit Enter once the Ubuntu welcome screen is displayed. This will start the installation program and get the ball rolling. However, by hitting the F2 key, and then the F1 key (the function keys run along the top of the keyboard), you can view a number of troubleshooting options. It's extremely unlikely you'll need to make use of these, but they're worth bearing in mind for future reference.

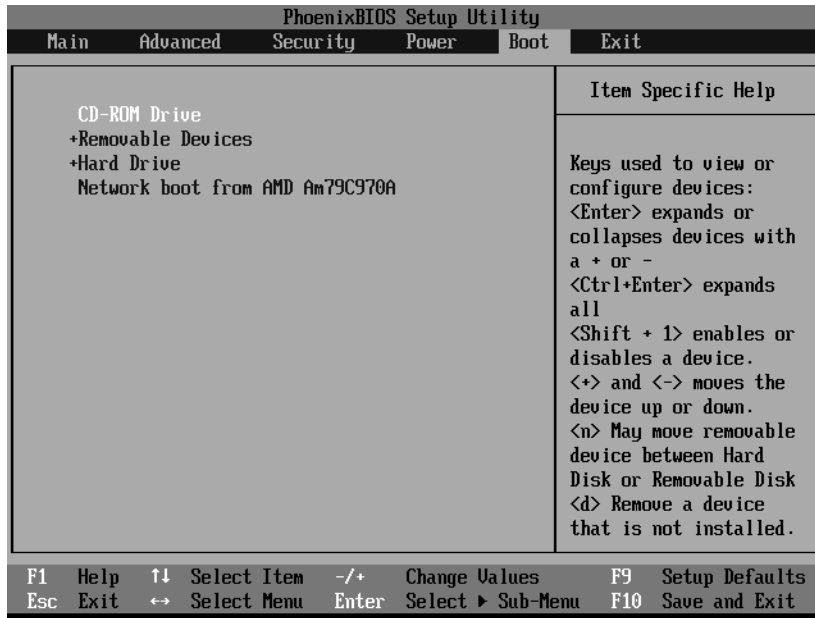


Figure 5-1. Before starting, make sure your computer can boot the CD-ROM.




Figure 5-2. Press Enter to start the installation program.

Step 2: Select Your Language

One of the design goals of Ubuntu is to be usable by just about anyone in the world. Ubuntu supports a massive list of languages, and the first step in the installation routine is to select one, as shown in Figure 5-3. It offers many eastern and western European languages, as well as Asian languages. The default is English.

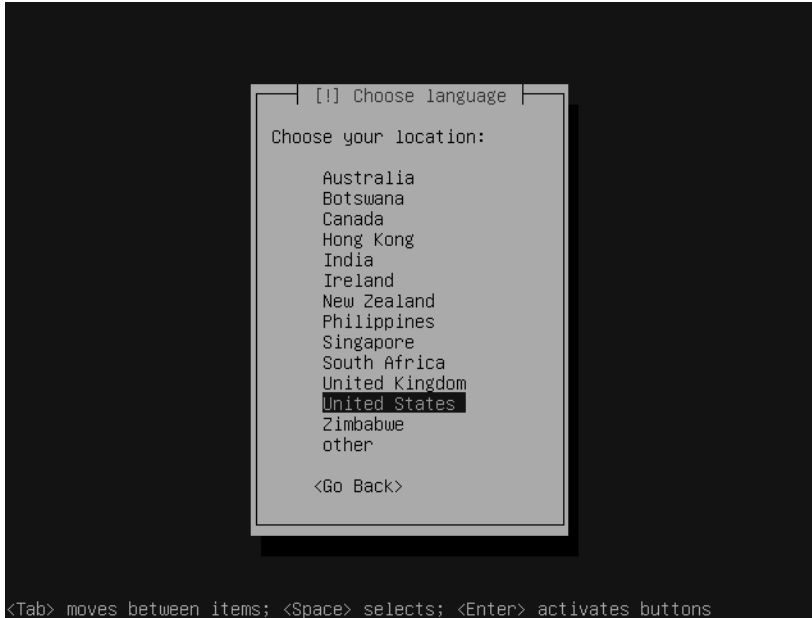


Figure 5-3. Choose your language and then your locale.

This will be your first experience with Ubuntu’s keyboard-based navigation system, and it’s a good example of how easy it is to use. Just use the up and down arrow keys to highlight whichever language you want to select, and then press the Enter key.

Following this, you may be asked to select your exact version of the language, which will be based on your locale. For example, residents of the United Kingdom might want to choose UK English. Once again, simply use the cursor keys to highlight your selection, and then hit Enter.

Step 3: Confirm Your Keyboard Layout

Next, you’ll be asked to confirm the keyboard layout you’ll be using. This should correspond to your language and locale settings, and will be automatically selected, so you can just hit Enter.

If, for whatever reason, you don’t know what keyboard layout you’re using, you can opt to discover it. Use the cursor keys to highlight “Find your layout by pressing some keys.” This will take you through a quiz during which the installation program uses a process of logical deduction to discover your keyboard layout.

Alternatively, if you think your choice is correct, you might want to test it by selecting the “Test your keyboard” option. If you find the keyboard doesn’t work as expected, choose the option to discover its layout.

Step 4: Wait During Hardware and Networking Autodiscovery

Next, you can sit back for a few moments while Ubuntu probes your hardware and configures everything to ensure that installation progresses smoothly. You'll see a progress bar indicating its status, as shown in Figure 5-4.

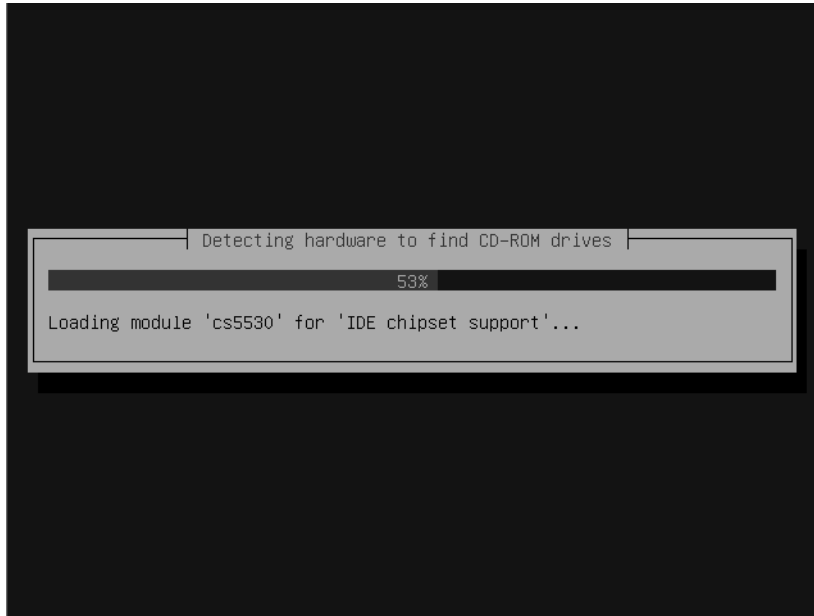


Figure 5-4. *Ubuntu will automatically discover your hardware to ensure the installation runs smoothly.*

Ubuntu will also attempt to configure your network connection at this stage. It may fail in this attempt, reporting an error relating to something called DHCP (Dynamic Host Configuration Protocol), but this isn't a problem. Just press Enter in the error dialog box, and then use the cursor keys to highlight "Do not configure the network at this time" (being offline won't hinder installation in any way).

Step 5: Set a Hostname

Following autodiscovery, you'll be asked to set a hostname for the computer, a slightly archaic requirement. The hostname is the name under which your computer is identified in networking environments, as well as at the command line. (Part 4 of this book describes how to use the command line.)

Although the hostname is essential to how Linux works, it's not something you need worry about. You can enter your own name or go with the default choice of ubuntu—it's up to you! Once you've made your choice, simply press Enter to move on.

Step 6: Partition the Disk

Partitioning the disk is one of the most important steps during installation although, unfortunately, it's one that's also couched in difficult terminology. Ubuntu does its best to make partitioning easy.

You have four main options when it comes to disk partitioning: install on a hard disk that also contains Windows, install on a second hard disk, delete an existing Windows partition, or use free space that already exists on your hard disk.

Sharing a Hard Disk with Windows

If you're thinking of installing Ubuntu on a hard disk that has Windows on it, Ubuntu can shrink the Windows partition to make space (provided you have enough free space; see Chapter 4 for details).

Simply select the top option in the menu, which should read something like "Resize IDE1 master, partition #1 (hda1) and use freed space," as shown in Figure 5-5 (what you see may vary slightly from this).

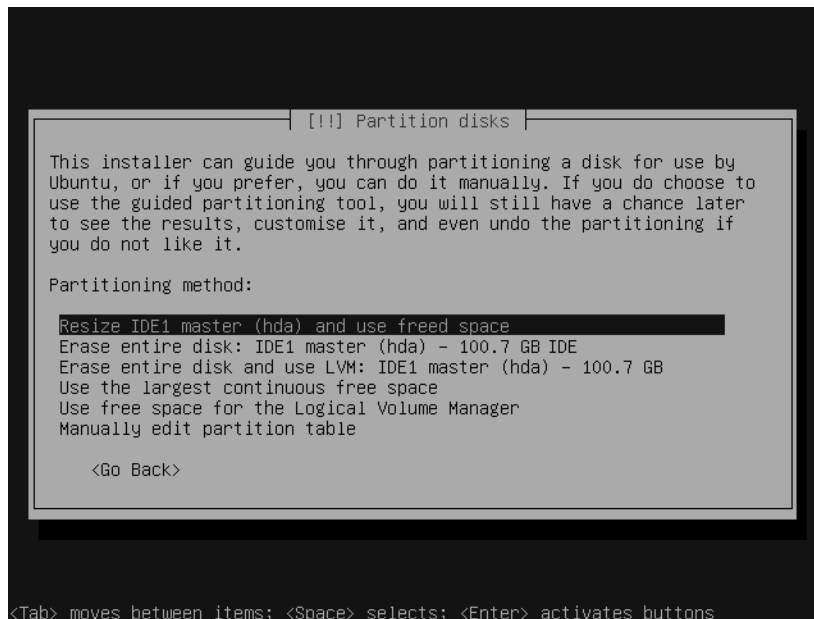


Figure 5-5. You can opt to shrink the Windows partition that's already on the disk.

Following this, you'll be prompted to type in a new size for the Windows partition. Ubuntu will tell you the minimum and maximum sizes for this partition. On my test PC, I was told the minimum size was 1.4GB and the maximum was 30GB. In other words, the Windows partition had around 1.4GB of data, and I couldn't shrink it beyond this.

You have two main considerations here: you must shrink the partition to have enough space for Ubuntu *but also* maintain enough space to allow Windows to function correctly. Although you might choose to shrink the Windows partition to its minimum size, bear in mind

that Windows needs a little spare space to write temp and system files, not to mention files for the users, such as Word documents. At the very least, you should allow around 500MB to 1GB on top of the minimum Windows partition size recommended by Ubuntu.

A handy hint is that you can simply type **50%**. This will split the free space in half, giving 50% to Ubuntu and leaving Windows with 50%. A neat and easy compromise!

Throughout the procedure, bear in mind that Ubuntu needs around 2GB to be able to function correctly, so you should create at least this much free space.

After the resizing, the automatic partitioning tool will take over. This automatically creates the partitions necessary for Ubuntu. You'll see a warning about writing changes to your disk, as well as a warning about how data may be destroyed. If you wish, you can highlight the Go Back option to check the choices, but in nearly all cases, the default choices made by Ubuntu are correct, so you can simply highlight Yes and hit Enter.

Deleting the Windows Partition

Deleting the Windows partition on your main disk will remove all the data from your hard disk, including the Windows operating system (if it's installed). Ubuntu will then take over the entire hard disk. If you're happy with this choice, simply use the cursor keys to select the "Erase entire disk: IDE1 master (hda)" choice.

Caution Be careful not to select the option marked "Erase entire disk and use LVM." This is an option for a different kind of Ubuntu setup used by experts.

Following this, you'll be asked to confirm the partitioning choices. The default options will be fine in most cases, so simply highlight Yes and hit Enter.

Installing on a Separate Hard Disk

Installing Ubuntu on a hard disk you've added especially for this purpose is very similar to installing it on the primary hard disk.

Beneath the menu option for erasing the disk should be a second similar option, along the lines of "Erase entire disk: IDE1 slave (hdb)," followed by the name and model of the second drive. You should select this option.

Caution Make sure that you select to install Ubuntu on the "slave" drive. *Do not select the option relating to the master drive!* Doing so will erase all the data on your primary disk, including your current Windows installation.

Once the choice has been made, the automatic partitioning tool will step in, and you can simply select Yes. Then press Enter to continue with the installation.

Using Existing Free Space

If you're an advanced user who has already freed space on your hard disk using a third-party partitioning tool, or if you deliberately set up Windows using a partition size that didn't take up the entire hard disk, you can simply select the "Use largest continuous free space" option.

Ubuntu will then automatically partition your hard disk, and all you need to do is confirm that you want to write the changes to disk.

Step 7: Wait During Automated Package Copying

The next stage is the automated copying from the CD-ROM of packages, as shown in Figure 5-6. Initially, these are copied to your hard disk. Later, they will be installed and set up.

Note *Packages* are single files that contain the individual program files and data. They're analogous to Windows program installation files.




Figure 5-6. *Ubuntu will copy all the necessary installation files to your hard disk.*

Unlike some varieties of Linux, Ubuntu doesn't ask you to make any choices regarding which software you want to install. It takes a "one size fits all" approach, which means that every installation of Ubuntu is exactly the same straight out of the box. However, you can add or remove programs very easily once the system is up and running.

Don't worry about the names of the packages as they flash by on the screen. Some might seem quite strange, complicated, or even irrelevant. The way Linux works means that some

software packages rely on other software packages, so often software that will never be used directly is installed.

Step 8: Set the Time Zone

Depending on your location, the next step may be to set the time zone for your locality, as shown in Figure 5-7. For example, users in eastern United States will choose Eastern Standard Time. This setting also allows Ubuntu to update your system clock to take into account Daylight Saving Time (or similar schemes used throughout the world).

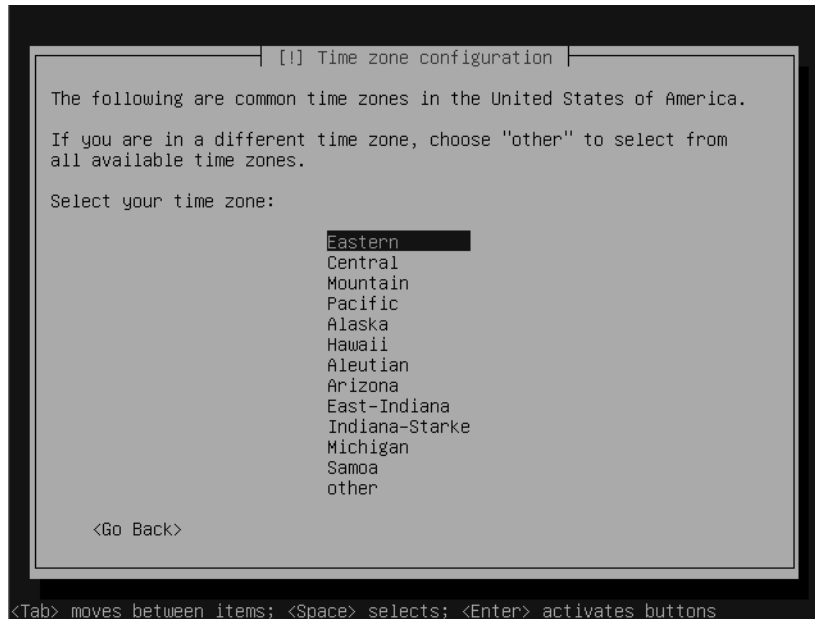


Figure 5-7. Choose the time zone appropriate to where you live.

Depending on your location, you may have an option to use GMT/UTC. If you're dual-booting with Windows, you should decline this option because it can cause your PC's clock to be set incorrectly. If Linux is the only operating system on your PC, you should choose GMT.

Step 9: Set Your Real Name, Username, and Password

Next, enter your real name and your username, as shown in Figure 5-8. The *real name* is how you'll be formally identified on the system to anyone who uses the system. The standard practice is to use your full name, including first and last names, separated by a space.

The *username* is how the computer itself will identify you. It's what you'll use to log in to the system, for example. This needs to be unique (two users on the same computer cannot have the same username), and there are rules about which characters you can use.

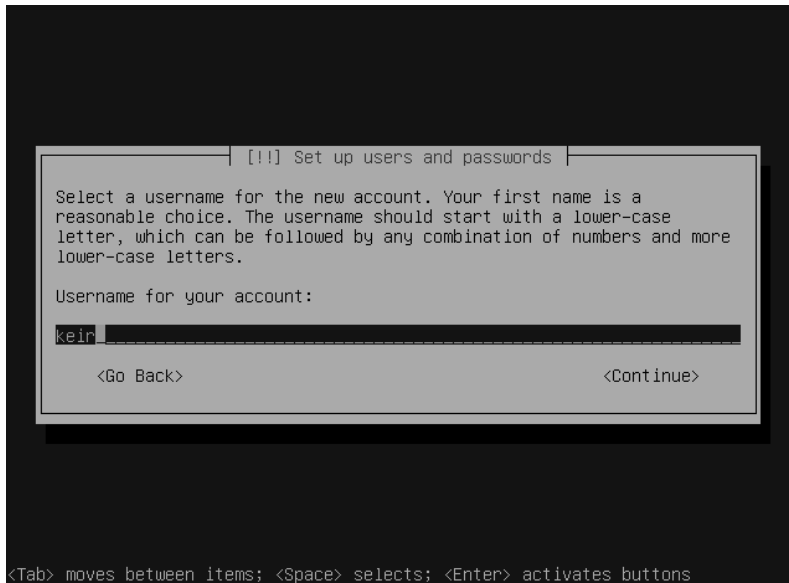


Figure 5-8. A good username is simple, and it's best to base it on your first name.

The username should be one word without any spaces in it. You can choose any username consisting of uppercase and lowercase letters, but try to avoid symbols and punctuation, because these can create problems later on. Additionally, the username cannot begin with an uppercase letter, although you can use uppercase in the rest of the username.

The simplest procedure for choosing a username is to use your own first name, typed entirely in lowercase letters. For example, when I installed Ubuntu, I typed my real name as Keir Thomas and chose keir as my username.

Following this, you'll be asked to enter a password. Here, the rules are the inverse of those for your username. A good password contains numbers, uppercase and lowercase letters, punctuation marks, and anything else you can get in there! This helps make your password almost impossible for someone else to guess, and thus makes your system more secure. (If you want to be *really* secure, create a password that's ten or more characters long.) You'll need to enter the password twice; the second time confirms that you didn't make a typo the first time around.

Step 10: Install the GRUB Boot Loader

If Windows is also installed on the same hard disk as Ubuntu, you'll be asked if you want to install the GRUB boot loader onto the PC's master boot record, as shown in Figure 5-9. The GRUB boot loader is the menu that will appear when you first boot, from which you'll be able to choose either Ubuntu or Windows. It's pretty essential, so you should select Yes. Then hit Enter.

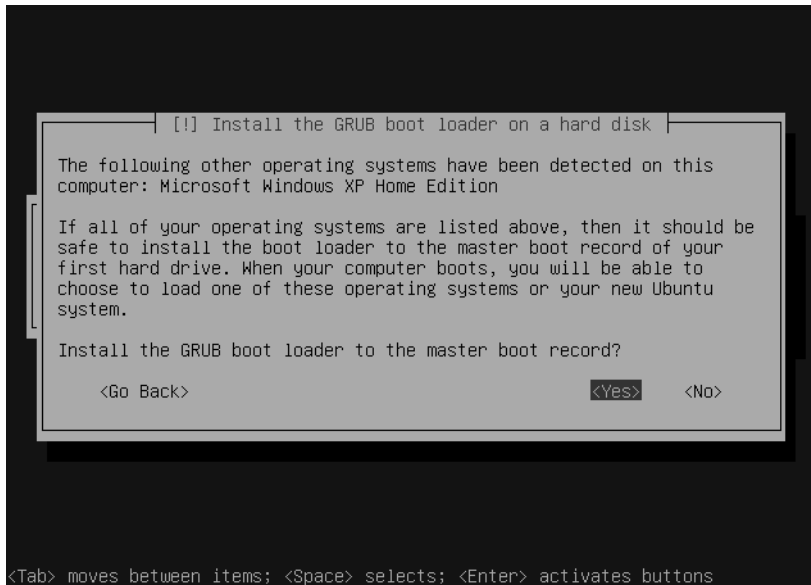


Figure 5-9. The GRUB boot loader is the menu that appears when you boot, allowing you to choose between Ubuntu and Windows.

Step 11: Reboot for the Second Installation Phase

Following a little more copying and installing, Ubuntu will reboot and begin the second phase of installation, as shown in Figure 5-10. You should make sure that the installation CD is no longer in the CD-ROM drive, and then press Enter. Note that you are now finished with the installation CD.



Figure 5-10. A quick reboot is needed to start the second installation phase. Make sure you remove the installation CD from the CD-ROM drive!

Step 12: Wait During Package Installation

Following the reboot of your PC, you will see lots of complicated-looking text scrolling down your screen. Don't worry—this is just the standard Linux boot procedure.

Soon after this, the graphical boot process will start, and you'll see the Ubuntu logo in the middle of the screen with a progress bar showing how far along it is. Beneath this a list of boot items will slowly appear as Ubuntu starts up. Again, don't worry if some of these read "Failed" instead of "OK." Error messages when Linux is booting aren't necessarily bad things, as they are with Windows! Often, they're just telling you something that you might want to take note of but that you can otherwise ignore.

After the boot process has finished, package installation will start. This will take anywhere between 15 and 30 minutes, during which time you'll see a progress bar on screen, as shown in Figure 5-11.




Figure 5-11. *Package installation will take up to half an hour.*

Step 13: Configure Your Monitor

Toward the end of the package installation, you may be asked to set the resolutions for your monitor, as shown in Figure 5-12, although this depends on your hardware setup. In many instances, Ubuntu will choose the correct settings and not bother you.

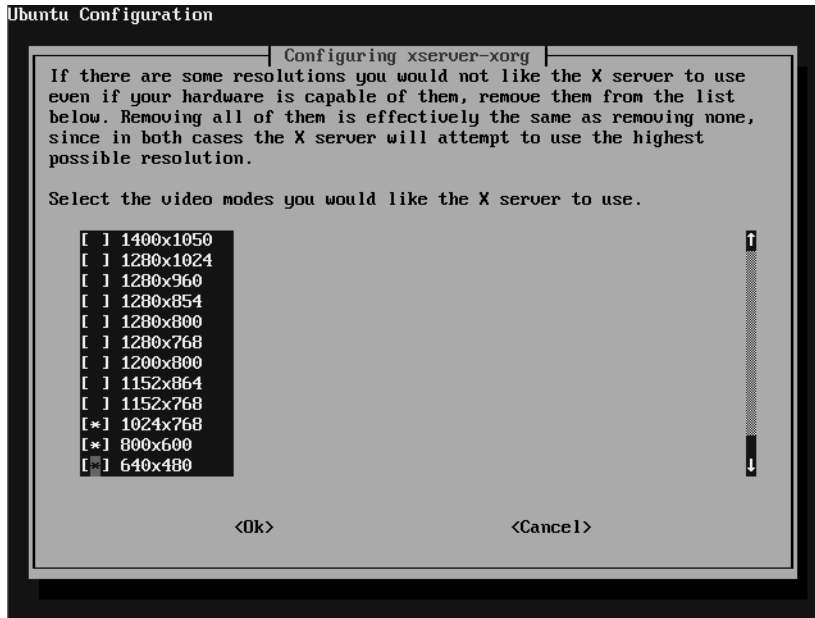


Figure 5-12. Choose your preferred monitor resolutions from the list.

If you are prompted for this information, you'll need to choose not only the main day-to-day resolution, but also any other resolutions that you might want to use occasionally, such as lower resolutions.

You should consult your monitor's manual to find out its optimal resolution, although Ubuntu will make a guess that's probably correct. Generally speaking, most 15- and 17-inch traditional CRT monitors run at 1024×768 resolution, while 17-inch LCD screens will most likely run at 1280×1024. If you have a 15-inch or 14-inch LCD screen, it will most likely run at 1024×768. Most modern notebook screens also tend to run at 1024×768.

Use the cursor keys to highlight the relevant option and, if it's not already selected (there isn't an asterisk there), press the spacebar to select it. Note that the list will scroll when you reach the bottom. When you're finished, use the Tab key to select OK, and then hit Enter.

Following this, package installation will continue for a bit longer.

Step 14: Boot for the First Time

When installation has finished, the system will boot straight into the Ubuntu login screen, as shown in Figure 5-13. From here, you can progress to Chapter 7 to learn how to get started. Alternatively, if you've run into any problems, see Chapter 6.



Figure 5-13. You'll know when installation is complete because the standard Ubuntu login screen will appear!

Summary

By following the steps outlined in this chapter, you should now have Ubuntu installed on your computer. I've tried to provide you with enough information to get around any problems, as well as explain exactly what's happening every step of the way.

Alas, it's still possible that you encountered hurdles that weren't addressed here. In the next chapter, you'll find solutions to common problems associated with Ubuntu installation.



Solving Installation Problems

Chances are that your installation of Ubuntu will complete without a hitch, and you'll find yourself with a first-rate operating system up and running within just a few minutes. However, if a problem rears its ugly head, you should be able to find the solution in this chapter, which addresses the most common installation problems. These problems are organized by when they occur: before you begin the installation, during the installation, and after the installation. The final section of the chapter describes how to configure the graphical subsystem with the X.org configuration utility, which can be useful if graphical glitches arise.

Preinstallation Problems

Some problems might arise before you even begin the installation process, or very early in the process, before the main installer program starts to run. This section addresses such issues.

Problem

My PC is unable to boot from the CD-ROM.

Solution

All modern computers are able to boot from a CD-ROM. Check your BIOS settings, as described in the previous chapter, to ensure that your BIOS is set correctly. In particular, check on-screen messages at boot time to see if any special key combination needs to be pressed to ensure your computer boots from the CD-ROM. On certain computers, you may need to press a key combination even if the BIOS is set to boot from the CD-ROM.

Some very old computers might not be able to boot from a CD-ROM. However, all computers without exception can boot from floppy disks, and Ubuntu includes files that can be written to a floppy disk in order to kick off the installation routine. To use this procedure, you'll need a blank 1.44MB floppy disk, the RawWrite program, and the Ubuntu installation CD-ROM. Follow these steps to write to the floppy disk and start the installation routine:

1. Boot into Windows and insert the Ubuntu CD-ROM into the drive.
2. Open your web browser and go to <http://uranus.it.swin.edu.au/~jn/linux/rawwrite.htm>. Select to download the latest version of RawWrite (click the link marked Binary).

3. Unzip RawWrite to an empty folder on the desktop and run the program.
4. Click the ellipsis button next to the box marked Image File. Browse to the \install folder on the Ubuntu CD-ROM. In the Files of Type drop-down list, select All Files. Then select `sbm.bin`.
5. Ensure the blank floppy disk is in the drive and click the Write button in the main Raw-Write window.
6. After a minute or two, the write process will be completed. Then simply reboot the machine, ensuring that both the floppy disk and Ubuntu CD-ROM are in the drives.
7. At the boot menu, use the arrow keys to move the selection down to CD-ROM. Then press Enter *twice* to start the Ubuntu installation routine.

Problem

The computer boots from the CD-ROM, but when the initial boot screen appears, hitting Enter doesn't start the installation. In fact, nothing happens at all!

Solution

If your PC uses a USB keyboard, it might be that it's not being recognized by the Ubuntu boot loader. It's possible to make most computers pretend that USB keyboards are older PS/2 keyboards. This is done on a fundamental hardware level and is invisible to the operating system. Here are the steps:

1. Enter the BIOS setup program by pressing Delete during the initial stages of your computer boot routine (while memory testing and drive identification are still taking place). Some computers might use a different key combination to enter BIOS setup, such as Ctrl+Insert, but this information will be displayed on screen.
2. Use the cursor keys to navigate to the Integrated Peripherals section, and then look for an entry along the lines of USB Legacy Support. Set it to Enabled.
3. Press Escape to return to the main menu and opt to save the changes.
4. Reboot the computer.

Note that you should repeat this procedure and deactivate USB Legacy Support once Ubuntu has been installed. At that stage, Ubuntu should be able to recognize the USB keyboard properly.

Problem

After I've hit Enter to clear the initial loading screen, the screen fills with text, and then the computer hangs with a message along the lines of "Kernel Panic."

Solution

Kernel Panic errors occur when the operating system cannot continue to load for various reasons. In this context, it's likely that either the CD is faulty (or dirty) or that your PC has a defective item of hardware.

First, check to make sure the CD is both clean and not scratched. If possible, try it on a different computer. If it works, then it's clearly not at fault, and your computer most likely has a hardware issue. In particular, bad memory can cause problems. Does the computer already have an operating system installed? Does this run without problems? If not, consider replacing your memory modules.

To thoroughly test your computer's memory, boot from the Ubuntu CD-ROM and type `memtest` at the initial prompt. Any problems with your memory will be reported in the Errors column on the right side of the program screen.

Problem

After I hit Enter on the loading screen, the screen fills with text. Then a message about configuring the framebuffer appears, but then the screen is corrupted and/or the text looks odd.

Solution

Your graphics card may be incompatible with the framebuffer mode used by Ubuntu's installation routine. You can overcome this problem by rebooting and, when you normally press Enter to start the installation routine, instead type the following, and then press Enter:

```
linux debian-installer/framebuffer=false
```

Problem

The installation program freezes or reboots almost as soon as it begins.

Solution

The power-saving feature in your computer is causing problems. Boot from the CD-ROM and, instead of pressing Enter at the boot prompt, type the following, and then press Enter:

```
linux acpi=off
```

Installation Problems

During installation, you may get error messages or experience other difficulties. This section offers some solutions to common installation problems.

Problem

Early in installation, I see an error message about the computer being unable to find a "DHCP address." Then I'm invited to configure the network manually.

Solution

The Ubuntu installer wants to go online during installation to download the very latest software. This isn't essential, however, and you can simply select "Do not configure network at this time" and press Enter. You can then update your software later when the operating system is up and running.

Problem

When the installation reaches the stage where it's configuring packages (after the reboot), it freezes on a particular package. It then returns a message along the lines of "An error was returned while trying to install."

Solution

The problem might be that the installation CD-ROM is faulty. It could be dirty, in which case you should consider wiping it with a tissue. Alternatively, it might be scratched, in which case you should seek a replacement. You might consider downloading an ISO file from www.ubuntulinux.com and burning your own replacement copy of the CD-ROM. This will require a CD-burning program such as Nero (www.nero.com).

Another possibility is that the installation program had a problem with your computer hardware. It has been reported that Ubuntu might have problems with some SATA hard disks and/or chipsets in recently manufactured computers. In such cases, it has also been reported that simply trying the installation again from scratch might fix the problem.

Problem

I'm using the same keyboard, mouse, and monitor across several computers courtesy of a keyboard, video, and monitor (KVM) switch. During installation, it appears that the autoprobng of my video hardware produces the wrong results.

Solution

It's impossible to correctly probe a monitor if a KVM switch is attached to the computer. Consider attaching the monitor directly to the computer for the duration of the installation.

Problem

During installation, everything freezes when the words "Configuring apt . . ." appear on the screen.

Solution

This problem appears to be caused by Ubuntu attempting to access online repositories during busy periods. Although the computer appears to have frozen, you might find that waiting 30 minutes or so causes the computer to start responding again, and it will then continue with the installation.

You might also try disabling your network connection for the duration of the installation. If you have a network card and/or modem attached to your computer, try removing it temporarily. Alternatively, if you have an Ethernet card, try unplugging the network cable.

Problem

The installation program freezes while it's copying across a file called `e2fsprogs-udeb`. (This same problem might occur when other packages are being copied.)

Solution

This problem may stem from an apparent incompatibility between Ubuntu and certain DVD-ROM drives based on a particular chipset (including certain Plextor drives). The solution is to install from an alternative CD/DVD-ROM drive if you have one. If not, try downloading the latest version of Ubuntu from www.ubuntulinux.com.

Postinstallation Problems

Problems might also occur after you install Ubuntu. This section addresses several possible postinstallation problems. This section covers only problems that appear immediately after installation—those that prevent Ubuntu from working correctly immediately after its first boot. Issues surrounding the configuration of hardware or software are dealt with in the next part of this book.

Problem

After installation has finished, I find myself with a blank screen (or a screen that's full of corrupted graphics). If I reboot, the computer appears to boot Ubuntu correctly but, again, it ends with a blank screen.

Solution

For some reason, your graphical configuration isn't working. See the "Graphical Problems" section later in this chapter for the steps to fix this problem.

Problem

During bootup, the Ubuntu logo and/or on-screen status messages are corrupted. However, the computer will eventually boot to the desktop just fine.

Solution

The Ubuntu graphical boot process is probably causing problems for your particular hardware configuration. Follow these steps to correct the problem:

1. At the Ubuntu desktop, select Applications ► Accessories ► Terminal.
2. At the command prompt, type the following:

```
sudo nano /boot/grub/menu.lst
```
3. Type your password when requested.
4. Your boot menu configuration file will open in a simple word processor. Use the cursor keys to navigate to the *first* line that reads as follows:

```
kernel /boot/vmlinuz. . .
```

This will be towards the end of the file and, unlike many of the lines above it, won't be preceded by a hash symbol.

5. At the end of the line, you will see the word `splash`. Move the cursor to that word and delete it.
6. Press `Ctrl+X` and `opt` to save the file.

Following this, the system's boot will be text only, although the system will boot into a graphical desktop as usual.

Problem

When I boot for the first time, the resolution is too low/high!

Solution

As soon as you're logged in, select `System ► Preferences ► Screen Resolutions` from the main menu (at the top of the screen). In the dialog box, choose a resolution more appropriate to your screen and a refresh rate to match. For most CRT screens, a refresh rate of 70 or 75 Hz should be fine; for LCD panels or notebooks, a refresh rate of 60 Hz is appropriate.

Problem

After booting up, my USB mouse and/or USB keyboard are not recognized!

Solution

Try unplugging the keyboard and/or mouse, and then reattaching them. If you find they now work, log in to Ubuntu, and then perform an online system upgrade. See Chapter 9 for more information.

If this fails to solve the problem, you can configure your BIOS to pretend your mouse and keyboard are traditional PS/2-style devices, as follows:

1. Enter the BIOS setup program by pressing `Delete` during the initial stages of your computer boot routine (while memory testing and drive identification are still taking place). Some computers might use a different key combination to enter BIOS setup, such as `Ctrl+Insert`, but this information will be displayed on screen.
2. Use the cursor keys to navigate to the `Integrated Peripherals` section, and then look for an entry along the lines of `USB Legacy Support`. Set it to `Enabled`.
3. Press `Escape` to return to the main menu and `opt` to save the changes.
4. Reboot the computer.

Problem

When I boot for the first time, all I see shortly after the PC is switched on is a command prompt that reads `GRUB>`.

Solution

For some reason, your installation of the GRUB boot menu went wrong. Try this solution, which involves reinstalling it:

1. Insert the Ubuntu installation CD-ROM and boot from it. At the splash screen (*before* the installation starts), type `rescue`.
2. The installer will start, and you'll need to answer the initial questions again, such as which language you wish to use.
3. Eventually, you'll be asked which partition you wish to mount (use for Ubuntu). Unless you know which partition contains your Ubuntu root partition, the only solution is to try each in turn. Select an option, and then enter the following command:

```
grub-install /dev/hda
```

4. If you receive an error message, start over from step 1 and try the next partition choice.
5. Remove the CD-ROM from the drive and reboot. You should now see the boot menu.

Problem

When I boot for the first time, all I see is a black screen with some text at the top saying, "Ubuntu 5.10 'Breezy Badger' Ubuntu tty1" and, beneath that, "ubuntu login:."

Solution

For some reason, the automatic configuration of your graphics card failed during installation. See the following section for instructions on configuring your GUI manually.

Graphical Problems

Although Ubuntu is extremely good at automatically detecting and configuring your PC's graphics hardware, it sometimes configures your GUI incorrectly. This is characterized by one of the following:

- Ubuntu freezes when the desktop would normally appear.
- You see on-screen graphical corruption, either of text or graphics.
- You see a message that the X server isn't working.
- You see a black screen with only with a text login prompt.

If you experience any of these problems, you will need to reconfigure Ubuntu's graphical subsystem, called X.org (often referred to simply as X). In most cases, you'll need your Ubuntu installation CD-ROM for this procedure.

To start the X.org configuration program, first boot into Ubuntu's recovery mode. This provides a simple command-line prompt and is designed to let you fix the system if anything goes wrong. If you dual-boot your computer with Windows, you'll need to select the "Ubuntu [...]" (recovery mode)" option from the menu that appears just after your computer boots. If your computer has only Ubuntu installed on it, you'll need to press a key to enter the boot menu when prompted. Then select the "Ubuntu [...]" (recovery mode)" option.

Eventually, a command prompt will appear, and you'll see `root@ubuntu:~#`, followed by a cursor. Type the following:

```
dpkg-reconfigure xserver-xorg
```

Reconfiguring X.org simply requires answering some simple questions. In many cases, default answers are supplied, and these should be sufficient. The configuration program guides you through selections for your graphics card, input devices, and monitor, as described in the following sections. Note that although the steps are separated into subsections for clarity, the program actually presents one long series of questions that moves seamlessly from topic to topic.

Graphics Card Configuration

The X.org configuration program begins with graphics card setup. Note that, like the installation program, this configuration program does not use the mouse; you make your selections using the keyboard, as described in Chapter 5.

1. The first option asks if you would like the configuration program to autodetect your video hardware, as shown in Figure 6-1. This is how Ubuntu originally attempted to configure your graphics card, and if you're having graphical problems, there's a good chance that this feature doesn't work correctly on your system. Therefore, you should select No.

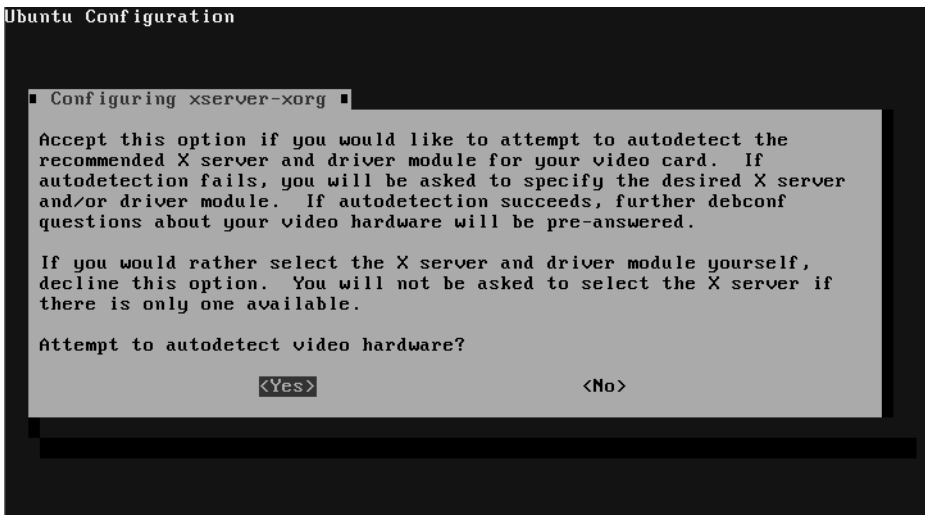


Figure 6-1. Select No to bypass autodetection.

2. You're asked to choose a driver from a list. If you wish, you can select the driver that seems appropriate to your hardware (the Ati option for an ATI-based card, for example), but the safest choice is to choose the Vesa option. This driver works on virtually every graphics card.

Note Using the Vesa driver entails a slight performance penalty, particularly when it comes to video playback. Therefore, you might choose to reconfigure X.org again in the future when your knowledge of Linux improves.

3. You're asked to enter an identifier for your graphics card. This is merely for reference purposes and anything will do. The default that Ubuntu suggests is fine.
4. An information screen outlining the next step, which is to identify the bus ID of the card, appears. This is a technical setting needed for Ubuntu to use the card. It should be automatically filled in, so highlight OK to confirm the default choice.
5. You're invited to enter the amount of memory your graphics card contains, as shown in Figure 6-2. Once again, Ubuntu is able to autodetect this. It's acceptable to leave the field blank and simply select OK.

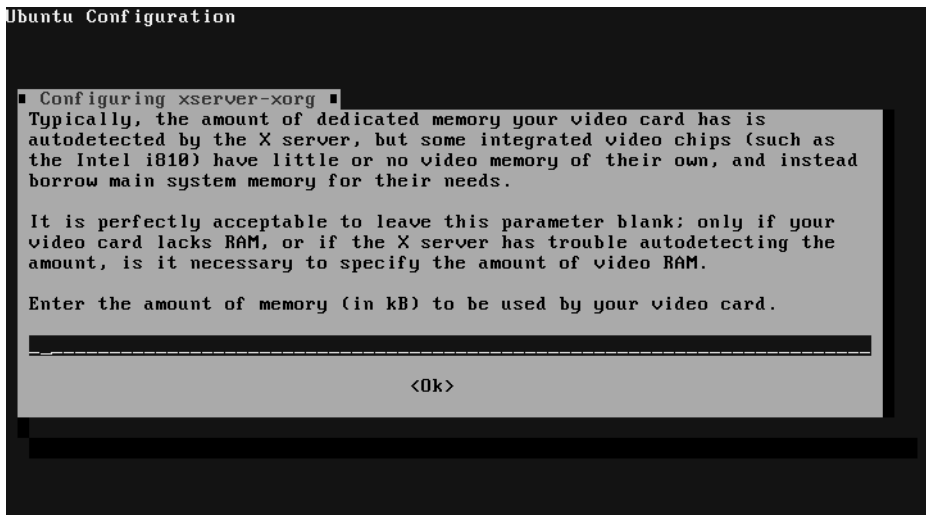


Figure 6-2. Ubuntu will automatically detect the quantity of graphics card memory, so you can leave the field blank.

6. You're asked whether you want to use your graphics card's framebuffer. This is a method of accessing the graphics card memory, and nearly all modern graphics cards are compatible with it, so you can select Yes. However, if after completing these steps you find that you still have problems with your display, start again from step 1 and choose not to use the framebuffer.

Keyboard and Mouse Configuration

Although your keyboard and mouse might be just fine, X.org handles the input devices as well as the graphics subsystem, and you now need to reconfigure them.

1. You're asked whether you want the configuration program to autodetect your keyboard layout. This is perfectly acceptable, so select Yes.
2. You're invited to enter your keyboard layout. This takes the form of two-letter country codes, the type used throughout the world. For example, a United States keyboard layout is indicated by typing US. United Kingdom users will want to type UK. Australian users should type AU. For a complete list of world country codes, see www.iso.org/iso/en/prods-services/iso3166ma/02iso-3166-code-lists/list-en1.html.
3. You're asked to enter the XKB rule set you wish to use. The default answer is fine for most users, so highlight OK.
4. You see an information screen describing keyboard models. In a nutshell, most recent English-based keyboards are pc104 models, which is to say they have 104 keys (see Figure 6-3). If you're using an older keyboard that does not have the Windows keys to the left and right of the spacebar, you should enter pc101. Highlight OK, press Enter, and then enter the correct choice.



Figure 6-3. *If your keyboard has a Windows key, it's a pc104 model. Otherwise, select pc101.*

5. Next is another information screen discussing the use of additional keyboard rules. If you have an English-based keyboard, you can ignore this step, and leave the field in the next screen blank.

6. You're asked if you wish to provide any special keyboard definitions. This is an interesting option you might like to explore in future, but you can ignore it for now. Simply leave the field blank on the next screen.
7. You're asked if you want to emulate a three-button mouse. Linux is designed around a three-button mouse system, rather than the two-button mouse used with Windows (or the single button used on the Apple Macintosh). Emulation of a three-button mouse means that when the left and right buttons are clicked at the same time, the computer thinks a third button has been pressed. Emulation isn't necessary with most modern mice because the scroll wheel is configured as the third button, so you should select No here.
8. You're asked if you want to configure the up and down action of a mouse scroll wheel as additional button presses. You can answer No to this option.
9. You're shown an information screen discussing modules. Select OK, and you'll be invited to select from a list. Once again, Ubuntu does the hard work for you, and the modules will be automatically selected. All you need to do is select OK at the bottom of the screen.
10. You're asked to confirm that you want to write the Files section to your X.org configuration file. You should certainly select Yes here.
11. You're asked to confirm that you want to add a DRI entry to your X.org configuration file. Again, select Yes. If you find your display still doesn't work correctly after completing these steps, you might start again, and select No for this option.

Monitor Configuration

Now it's time to configure your monitor.

1. As before, you can choose to autodetect the monitor, but if your graphical system isn't working, there's a chance that this feature is incompatible with your system. Therefore, you should opt not to autodetect.
2. You're asked to enter an identifier for your monitor. Anything will do, and the default Ubuntu suggests is usually fine.
3. Select from the list the resolutions you would like to be available to use under Ubuntu. Highlight each entry in the list and press the spacebar, so that an asterisk appears alongside it. See Table 6-1 for guidance on which resolutions you should choose. Most notebooks run at 1024×768, regardless of screen size, although some more high-end models might run at higher resolutions, particularly if they are wide-screen models (refer to your notebook documentation).

Note TFT screens are designed to run at a single optimal resolution, rather than at a range of resolutions. Therefore, you should select only one resolution from the list if you have this type of monitor.

4. You're asked to enter some technical characteristics of your monitor. You can choose the Simple, Medium, or Advanced option; each requires you to enter progressively more information. In most cases, the Simple option should suffice and simply involves entering the physical size of your screen (such as 15 or 17 inches). However, if after completing these steps you find you still have problems, you might try repeating the configuration and choosing the Advanced option, which will ask for the horizontal and vertical scan rates of the monitor. You can find this information in the monitor's manual or, if that's not available, by searching online for the specification list of your model, as shown in Figure 6-4.

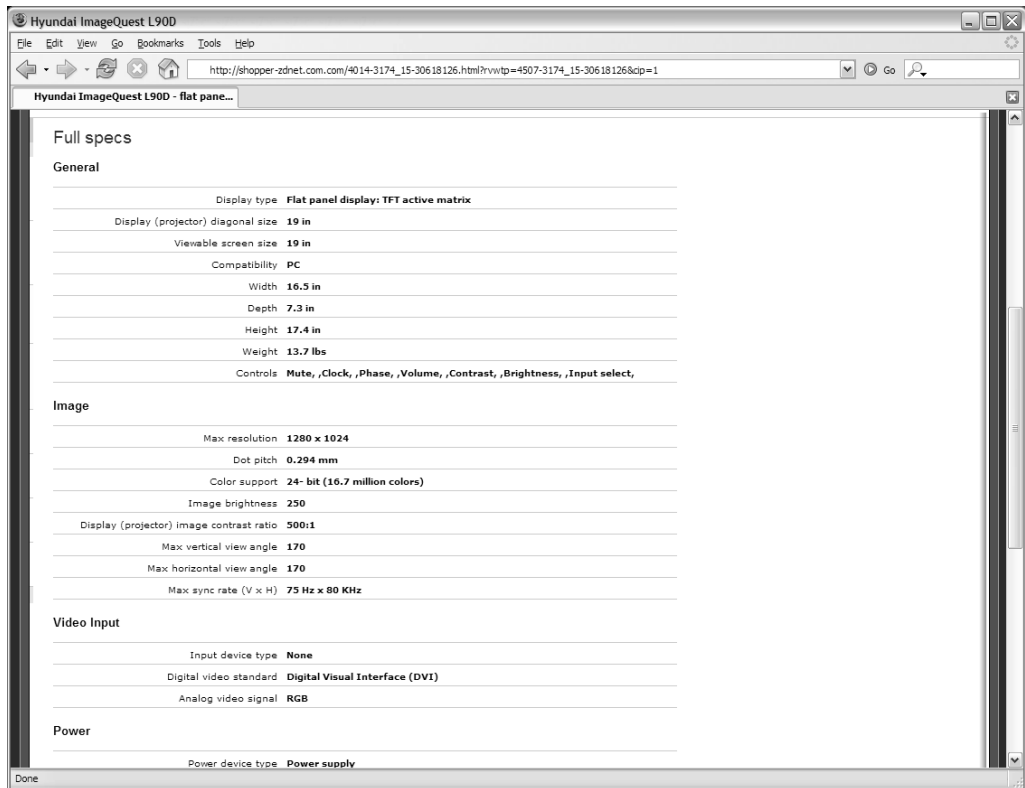


Figure 6-4. A quick search on the Web will reveal the specifications of most monitors.

5. You're asked to confirm that you want to write the monitor synchronization ranges to the configuration file. You can select Yes here.
6. You're asked to enter the color depth you wish to use. All modern graphics cards are capable of 24-bit color, but if you have an older card with less memory, you might choose 16 from the list. Once again, if you find that the graphical configuration doesn't work after completing these steps, you might wish to return to this step and try decreasing the value.

7. After selecting the color depth, the configuration file will be written to disk, and you'll be returned to the command prompt. You can now reboot and try out the new settings.

Table 6-1. *Typical Monitor Resolutions*

Monitor Size	Typical Resolutions
CRT Monitors	
14 inches	800×600, 640×480
15 inches	800×600, 640×480
17 inches	1024×768, 800×600, 640×480
19 inches	1280×1024, 1024×768, 800×600, 640×480
20 inches	1600×1200, 1280×1024, 1024×768, 800×600, 640×480
TFT Screens	
14 inches	1024×768
15 inches	1024×768
17 inches	1280×1024
19 inches	1280×1024

Summary

This chapter's goal was to address problems that might occur during the installation of Ubuntu. It discussed preinstallation, installation, and postinstallation issues. It also covered how to use the X.org reconfiguration utility to configure the graphics subsystem, which may be necessary if the installation program failed to properly recognize your graphics card or monitor.

You should now have Ubuntu installed. The next part of this book focuses on helping you get everything up and running. You'll learn essential skills and become a confident Linux user.

PART 3



The No-Nonsense Getting Started Guide



Booting Ubuntu for the First Time

Now that Ubuntu is installed, you'll no doubt want to get started immediately, and that's what Part 3 of this book is all about. In later chapters, we'll examine specific details of using Ubuntu and getting essential hardware up and running. We'll also look at personalizing Ubuntu so that it works in a way that's best for you on a day-to-day basis. But right now, the goal of this chapter is to get you doing the same things you did under Windows as quickly as possible.

This chapter explains how to start up Ubuntu for the first time and work with the desktop. It also looks at how some familiar aspects of your computer, such as using the mouse, are slightly enhanced under Ubuntu.

Starting Up

If you've chosen to dual-boot with Windows, the first Ubuntu screen you'll see is the boot loader menu, which appears shortly after you switch on your PC. If Ubuntu is the only operating system on your hard disk, you'll see a brief one-line message pointing out that if you press a key, you can access this boot menu. You won't need to do so unless you want to access the recovery mode boot settings. In fact, if Ubuntu is the only operating system on your computer, you can skip to the next section.

Note The boot loader is actually a separate program called GRUB. This program kicks everything off and starts Ubuntu.

The boot loader menu you see when your PC is set to dual-boot has three or four choices, as shown in Figure 7-1. The top one is what you need to boot Ubuntu. The Ubuntu option will be selected automatically within ten seconds, but you can press Enter to start immediately.

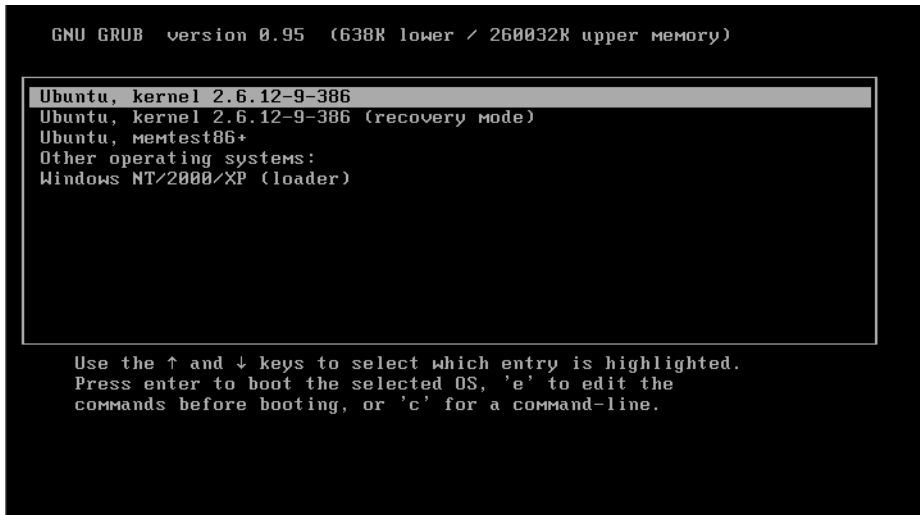


Figure 7-1. *The default choice is fine on the boot menu, so press Enter to start Ubuntu.*

You should find that you also have an entry for Windows, located at the bottom of the list. To boot into Windows, simply use the cursor keys to move the selection to the appropriate option, and then press Enter.

As well as entries for Ubuntu and Windows on the boot list, you should see an entry ending in “(recovery mode).” This is a little like Safe Mode within Windows, in that the system boots with conservative system settings, and you should be able to make repairs to the system if something has gone wrong. However, in Ubuntu’s recovery mode, you’re presented with a command-line prompt.

When you update your system software, you might find new entries are added to the boot menu list. Without exception, the topmost entry is the one you’ll want each time to boot Ubuntu. This will ensure you start the system with the newest version of the system software. The entries beneath this will start the system with older versions of the system software and are provided in the unlikely situation that the latest software causes problems.

Note All operating systems need a boot loader—even Windows. However, the Windows boot loader is hidden and simply starts the operating system. Under Ubuntu, the boot loader usually has a menu, so you can select Linux or perhaps an option that lets you access your PC for troubleshooting problems. When you gain some experience with Ubuntu, you might choose to install two or more versions of Linux on the same hard disk, and you’ll be able to select among them using the boot menu.

Logging In

After Ubuntu has booted, you should see the login screen, as shown in Figure 7-2. Here, you enter the username and the password you created during the installation process. By clicking

one of the buttons on the bottom-left side of the screen, you can also opt to reboot the system or shut it down.



Figure 7-2. Type your username, then enter your password, and then press Enter to log in.

The user account you created during installation is similar to what Windows refers to as an “administrator” account. This means that the account you use on a day-to-day basis can also change important system settings and reconfigure the system. However, the main difference between Ubuntu and Windows is that you’ll need to enter your password to make any serious changes. Don’t worry about damaging anything accidentally; trying to reconfigure the system or access a serious system setting will invariably bring up a password prompt. You can simply click the Cancel button if you don’t want to continue.

Note Unlike some versions of Linux, Ubuntu doesn’t encourage the user to use an actual root (or administrator) account. Instead, it operates on the principle of certain ordinary users having superuser privileges that allow them to administer the system. The user account you create during setup has these privileges.

Exploring the Desktop

After you've logged in, you'll see the welcoming tan and brown theme of the Ubuntu desktop, as shown in Figure 7-3. Feel free to click around and see what you can discover. There's little chance of you doing serious damage, so let yourself go wild and play around with your new operating system!

Tip Although you can't damage the system by messing around, you might find that you somehow cause programs to work incorrectly. Don't worry if this happens. You can always create a new account for yourself following the instructions in Chapter 30. When using this new account, you should find all the settings are returned to normal, and you'll be back to square one!



Figure 7-3. Feel free to experiment with the Ubuntu desktop and see what you can discover.

First Impressions

The first thing you'll notice is that the desktop is clean compared to Windows. You don't have a lot of icons littering the screen.

Of course, you can fill the desktop with stuff if you want to do that. As with Windows, you can save files to the desktop for easy access. In addition, you can click and drag icons from any of the menus onto the desktop in order to create shortcuts.

Along the top of the desktop, you see three menus:

- The Applications menu is the equivalent of the Windows Start menu. Here, you'll find access to all the software available under Ubuntu.
- The Places menu is somewhat like My Computer in Windows, in that it gives quick access to locations within the file system.
- The System menu is a little like the Windows Control Panel, in that it allows you to change various system settings.

The counterpart of the Windows Recycle Bin lives at the bottom right of the screen as a small icon and is called the Wastebasket. Although diminished in stature compared to the Windows representation, it works in a similar way: you can drag icons and files onto the icon in order to delete them, and you can click it to open the Wastebasket and salvage files.

The mouse works largely as it does in Windows, in that you can move it around and click on things. You can also right-click virtually everything and everywhere to bring up context menus, which usually let you alter settings. And you should find that the scroll wheel in between the mouse buttons lets you scroll windows.

Whenever Ubuntu is busy, an animated circular icon will appear that is similar in principle to the hourglass icon used in Windows. It also appears when programs are being launched.

You can shut down or reboot your PC by selecting the System ► Logout menu option.

Caution Bear in mind that Ubuntu isn't a clone of Windows and doesn't try to be. Although it works in a similar way—by providing menus, icons, and containing programs within windows—there are various potholes in the road that can trip up the unwary.

WRONG RESOLUTION!

You might find when you boot up that Ubuntu has defaulted to the wrong resolution. In other words, everything might be a little too large or too small. You might have trouble reading text, for example, or you might find that program windows fill the screen to the extent that their contents partially disappear off the edges.

Changing the resolution is simple. Select System ► Preferences ► Screen Resolution from the menu (at the top of the screen). In the Resolutions drop-down list, select the appropriate resolution for your monitor. For a 17-inch CRT monitor, the standard resolution is 1024×768 (although some people prefer 800×600). Most 17-inch TFT screens run at 1280×1024 resolution. If you have a 15-inch CRT monitor (common on PCs made before 2000), you'll probably find 800×600 a maximum setting; others prefer 640×480. A 15-inch TFT screen will usually run at 1024×768 resolution. If in doubt as to your monitor's resolution, consult your monitor's manual for more information.

Desktop Elements

The Ubuntu desktop is similar to that of Windows. It has the following elements:

- **Menus:** The three menus at the top left of the screen provide access to all of Ubuntu's functionality. As noted earlier, the Applications menu provides access to programs, the Places menu provides access to the file system, and the System menu provides access to configuration settings (as well as the Log Out option). You can click and drag practically every menu entry onto the desktop in order to create a quick-and-easy shortcut.
- **Icons:** Although the Ubuntu desktop is clean, some icons are tucked away at the top and bottom of the screen. Those at the top are located to the right of the menus and allow you to start the browser, e-mail client, and help system (and are arranged in that order). At the top right is a speaker icon that lets you alter the sound volume, while at the bottom left is the Hide Windows button that instantly minimizes all open windows to give access to the desktop underneath. At the bottom right are the four virtual desktop buttons, which I'll discuss in the "Working with Virtual Desktops" section later in this chapter, and also the Wastebasket icon.

Note If you're dual-booting with Windows, you might see an icon at the top left that will let you access your Windows files. Similarly, if you have a memory card reader or digital camera plugged into your PC, you might see desktop icons for them, too.

- **Panel:** The bar at the bottom of the screen, called the Panel, shows which programs are currently running (if any). As with Windows, you can simply click the button for any program to bring that window "to the top." Alternatively, you can right-click each entry to instantly minimize or maximize that particular window.
- **Clock:** The clock is located in the top right of the screen. Clicking it brings up a handy monthly calendar. Right-clicking it brings up a context menu. On this menu, the Preferences option lets you alter the way the date and time are displayed. The Adjust Date & Time option lets you change the time and/or date if they're incorrect.
- **Notification area:** The speaker icon and clock are located in the notification area, which is similar to the Windows system tray. Programs that like to hang around in memory, such as the Rhythmbox media player, will add an icon in this top-right area, to allow quick access to their functions. Usually, you simply need to click (or right-click) their icons to access the program features.

Tip The small bar marks the leftmost boundary of the notification area. To resize the notification area, right-click this bar and remove the tick from the Lock to Panel menu entry. Then you can click and drag the bar to a different size. This might be handy if the notification area starts to fill up with icons!

BEHIND THE DESKTOP: GNOME

Although I refer to the Ubuntu desktop, the fundamental software behind it is created by GNOME: The Free Software Desktop Project. This is one of the most well-established organizations currently producing desktop interfaces for Linux, as well as for other versions of Unix. Its home page is www.gnome.org.

Although it's based on GNOME, Ubuntu's desktop has its own set of individual features and programs, as well as a unique look and feel. That said, it works in an almost identical way to versions of GNOME that are used in other Linux distributions, such as Red Hat Fedora.

The nature of open-source software—whereby anyone can take the source code and create his or her own version of a program—makes Ubuntu's remodeling of the GNOME desktop possible. Unlike with Windows software, more than one current version of a particular program or software suite can exist, and each is usually tailored to the particular needs of one of the various Linux distributions.

Quick Desktop Guides

Refer to Figure 7-4 for an annotated diagram of the desktop. The figure includes an open menu, browser window, and program window, so you can get an idea of working from the desktop.

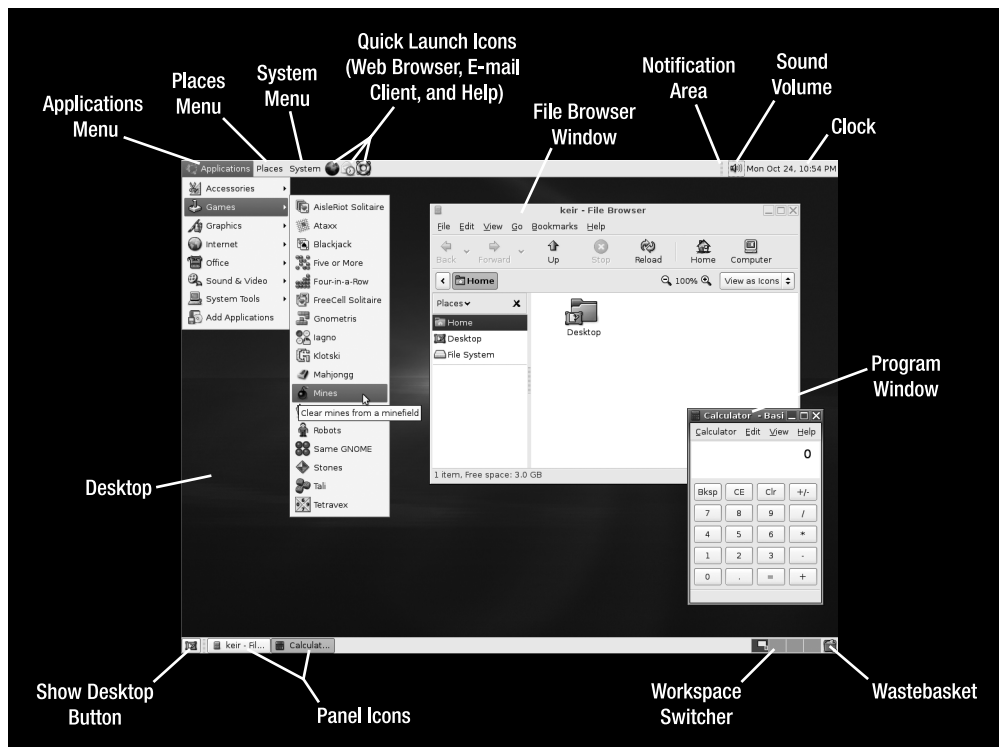


Figure 7-4. The Ubuntu desktop is broadly similar to the Windows desktop, with a few minor differences.

As another handy reference, Table 7-1 lists standard Windows desktop features and where similar functionality can be found on the Ubuntu desktop.

Table 7-1. *Windows Desktop Features Equivalents Under Ubuntu*

Windows Function	Description	Ubuntu Equivalent
My Computer	Double-clicking the My Computer icon gives you access to the PC system. In particular, it lets you browse the file system.	Click Places ► Computer to see all the drives attached to the computer in the file browser window. If you wish to browse the file system, double-click File System in the list on the left side of the file browser window.
Recycle Bin	The Recycle Bin is the repository of deleted files.	Click the small Wastebasket icon, located at the bottom-right corner of the Ubuntu desktop.
Start menu	The Start menu provides access to many computer functions, as well as a list of the programs installed on the system.	This function is split between the Applications and System menu. The Applications menu provides access to every piece of software installed under Ubuntu. The System menu lets you configure and administer the system, rather like the Windows Control Panel.
Quick Launch toolbar	Located just to the right of the Start button, these small icons let you launch popular programs with a single click.	Similar icons are located to the right of the main menus at the top of the Ubuntu desktop. You can add your own entries here by clicking and dragging program icons from the Applications menu.
My Network Places/ Network Neighborhood	This icon is used to access network services, usually within a business environment (on newer versions of Windows, this icon is often hidden by default).	To browse the local network, click Places ► Network Servers.
My Documents	The My Documents folder, accessed via its icon on the Windows desktop, is a storage space set aside for a user's documents.	The user's Home folder serves this purpose and can be accessed by clicking Places ► Home Folder.
Control Panel	The Windows Control Panel, located off the Start menu, allows the user to change system settings and preferences.	Similar functionality can be found under the System ► Administration and System ► Preferences menu options.
Find Files	Located on the Start menu, the Find Files function lets a user search the file system for missing items.	To find files, click Places ► Search for Files.

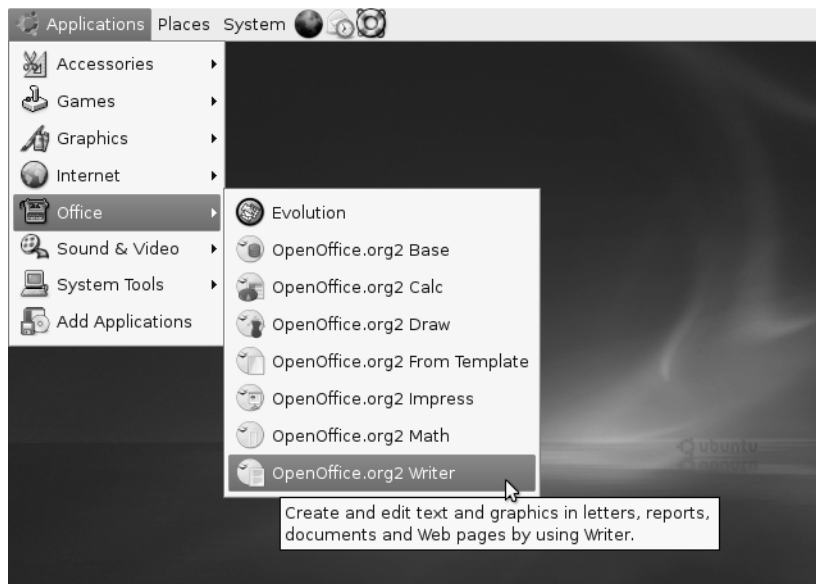
Table 7-1. *Windows Desktop Features Equivalents Under Ubuntu*

Windows Function	Description	Ubuntu Equivalent
Shutdown/Reboot	At the bottom of the Start menu within Windows is the Shutdown button.	Clicking System ► Log Out brings up a dialog box that is almost identical to the one displayed in Windows. It gives you the option of rebooting, shutting down, logging out (returning to the initial username/password prompt), or hibernating the system (suspending the RAM to disk and then shutting down the computer).

It will take some time to get used to the look and feel of Ubuntu; everything will initially seem odd. You'll find that the on-screen fonts look a little different from those in Windows, for example. The icons also won't be the same as you're used to in Windows. This can be a little disconcerting, but that feeling will quickly pass, and everything will become second nature. We'll look at how you can personalize the desktop in Chapter 10.

Running Programs

Starting a new program is easy. Just click the Applications menu and then choose a program from the list, as you would in Windows using the Start menu. The menu, shown in Figure 7-5, is split into various subcategories of programs, such as office tools, graphics programs, and even games!

**Figure 7-5.** *The programs on the Applications menu are split into various categories.*

If you want to start the web browser or e-mail client (arguably two of the most popular programs offered by Ubuntu), you can click their icons on the desktop, just to the right of the menus at the top of the screen (see Figure 7-4).

At the top right of every program window under Ubuntu, you'll see the familiar Close, Minimize, and Maximize buttons, albeit with a slightly different look and feel than you're used to. Clicking Close will end each program, as in Windows.

Working with Virtual Desktops

Windows works on the premise of everything taking place on top of a “desktop.” When you start a new program, it runs on top of the desktop, effectively covering up the desktop. In fact, all programs are run on this desktop, so it can get a bit confusing when you have more than a couple of programs running at the same time. Which Microsoft Word window contains the document you're working on, rather than the one you've opened to take notes from? Where is that My Computer window you were using to copy files?

Ubuntu overcomes this problem by having more than one desktop area. By using the Workspace Switcher tool, located at the bottom right of the desktop, you can switch between four virtual desktops. This is best explained by a demonstration.

1. Make sure that you're currently on the first virtual desktop (click the leftmost square on the Workspace Switcher) and start up the web browser by clicking its icon at the top of the screen (the globe icon located to the right of the menus).
2. Click the second square on the Workspace Switcher. This will switch you to a clean desktop, where no programs are running—desktop number two.
3. Start up the file browser by selecting the Places ► Home menu option. A file browser window appears.
4. Click the *first square* in the Workspace Switcher again. You should switch back to the desktop that is running the web browser.
5. Click the *second square*, and you'll switch back to the other desktop, which is running the file browser.

Tip Right-clicking any of the program entries in the Panel will bring up a menu where you can move a program from one virtual desktop to another. Just select Move to Another Workspace.

See how it works? You can create more than a couple virtual desktops—as many as 36! To set the number of workspaces, right-click the Workspace Switcher and select Preferences. In the window that appears, click the up arrow next to the Number of Workspaces entry, as shown in Figure 7-6. You might choose to have fewer than four virtual desktops. For example, some people use only two: one main desktop and one “spare,” to take the excess when the main desktop gets full.



Figure 7-6. Four virtual desktops are set up by default, but you can have as many as 36.

You can also rename each virtual desktop by double-clicking its entry in the list in the Preferences window. This allows you to be even more organized. For example, you might reserve desktop 1 for running Internet programs and give it a name that indicates this, such as Net Programs. You might then use desktop 2 to run office programs, giving it an appropriate title; use desktop 3 for file browsing; and so on. This name will then appear whenever you right-click a program's entry on the Panel and attempt to move it to a different desktop. Additionally, these titles will appear whenever you right-click and choose to send each program window to a different desktop.

Tip Putting your mouse over the Workspace Switcher and scrolling the mouse wheel switches between the various virtual desktops instantly.

The Workspace Switcher provides a way of organizing your programs and also reducing the clutter. You can experiment with virtual desktops to see if you want to organize your work this way. Some people swear by them. Experienced Ubuntu users may have in excess of ten virtual desktops, although clearly this will appeal only to organizational geniuses! Other users think multiple desktops are a waste of time. They're certainly worth trying out to see if they suit the way you work.

Using the Mouse

As noted earlier, the mouse works mostly the same under Ubuntu as it does under Windows: a left-click selects things, and a right-click usually brings up a context menu. Try right-clicking various items, such as icons on the desktop or even the desktop itself.

Tip Right-clicking a blank spot on the desktop and selecting Create Launcher lets you create shortcuts to applications. Clicking Create Folder lets you create new empty folders.

You can use the mouse to drag icons on top of other icons. For example, you can drag a file onto a program icon in order to run it. You can also click and drag in certain areas to create an elastic band and, as in Windows, this lets you select more than one icon at once.

You can resize windows using the mouse in much the same way as in Windows. Just click and drag the edges and corner of the windows. In addition, you can double-click the title bar to maximize and subsequently restore windows.

Ubuntu also makes use of the third mouse button for middle-clicking. You might not think your mouse has one of these but, actually, if it's relatively modern, it probably does. Such mice have a scroll wheel between the buttons, and this can act as a third button when pressed.

In Ubuntu, the main use of the middle mouse button is in copying and pasting, as described in the next section. Middle-clicking also has a handful of other functions; for example, middle-clicking the title bar of any open window will switch to the window underneath.

Tip If your mouse doesn't have a scroll wheel, or if it has one that doesn't click, you can still middle-click. Simply press the left and right mouse buttons at the same time. This emulates a middle-click, although it takes a little skill to get right. Generally speaking, you need to press one button a fraction of a second before you press the other button.

Cutting and Pasting Text

Ubuntu offers two separate methods of cutting and pasting text. The first method is identical to that under Windows. In a word processor or another application that deals with text, you can click and drag the mouse to highlight text, right-click anywhere on it, and then select to copy or cut the text. In many programs, you can also use the keyboard shortcuts of Ctrl+X to cut, Ctrl+C to copy, and Ctrl+V to paste.

However, there's a quicker method of copying and pasting. Simply click and drag to highlight some text, and then immediately click the middle mouse button where you want the text to appear. This will copy and paste the highlighted text automatically, as shown in Figure 7-7.

This special method of cutting and pasting bypasses the usual clipboard, so you should find that any text you've copied or cut previously should still be there. The downside is that it doesn't work across all applications within Ubuntu, although it does work with the majority of them.

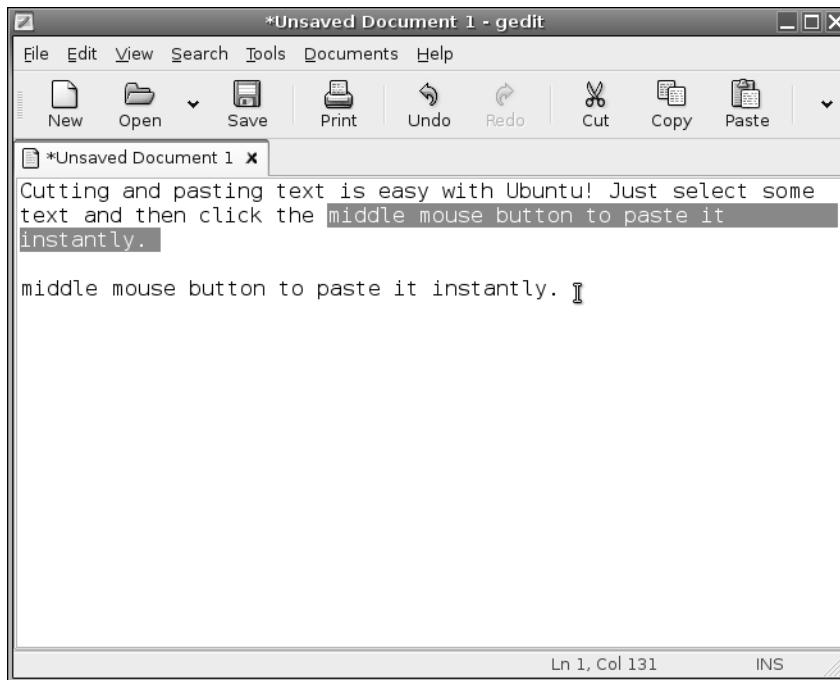


Figure 7-7. Highlight the text, and then middle-click to paste it instantly.

Summary

This chapter covered booting into Ubuntu for the first time and discovering the desktop. We've looked at starting programs, working with virtual desktops, using the mouse on the Ubuntu desktop, and much more. You should have become confident in some basic Ubuntu skills and now be ready to learn more!

In the next chapter, we'll look at getting your system up and running, focusing in particular on items of hardware that experience day-to-day use.



Getting Everything Up and Running

This chapter guides you through setting up all the essential components of your Ubuntu installation. This includes hardware configuration, as well as setting up e-mail and online software repositories. It covers the post-installation steps necessary for getting your system up and running efficiently.

Like all modern Linux distributions, Ubuntu is practically automated when it comes to setting up key hardware and software components. Key software will work from the start, and hardware will be automatically configured. However, you might need to tweak a few settings to make everything work correctly. Read on to learn more.

Ubuntu Hardware Support

The age-old criticism that the Linux operating system lags behind Windows in terms of hardware support is long dead. The majority of add-ins, such as digital cameras and printers, will all work with Ubuntu immediately, with little, if any, configuration.

In fact, most underlying PC hardware is preconfigured during installation without your knowledge and without further work required. Both your graphics card and sound card should work without a hitch, for example. In addition, nearly all USB and FireWire devices you plug in after installation will be supported.

Ubuntu doesn't support a few hardware items. Generally, it's a black or white situation; Ubuntu either works with a piece of hardware or it doesn't.

The types of hardware that Ubuntu doesn't support tend to be esoteric devices that rely on custom software provided by the hardware manufacturer. It's also sometimes the case that brand-new models of hardware won't work with Ubuntu because support has yet to be added. However, as soon as a new piece of hardware comes out, work is usually undertaken to ensure that Linux is made compatible with it. This work is carried out by the people who produce the Linux kernel, although sometimes the companies behind the hardware will supply their own code. This means that you can get the latest hardware drivers by grabbing the latest kernel files. This is one more reason why you should regularly update your system online, as explained in Chapter 9.

Tip Before you buy a new piece of hardware, why not ask the salesperson if it runs under Linux? You can only hope that the salesperson knows or can find out for you. Also, compatibility with Linux is often listed on the hardware box or at the manufacturer's web site (even if you sometimes need to search through the FAQ section to find out about it!).

Unfortunately, unlike with Windows, it's rare to find Linux drivers on the CD that comes with the hardware. Even if you do find a Linux driver supplied, chances are that it will work with only certain versions of Linux, such as Red Hat or SUSE Linux. At the time of writing, Ubuntu has yet to gain the kind of momentum where manufacturers specifically produce drivers for it. But this may change in the future.

Note It's possible to use a program called `alien` to convert packages designed for other distributions into Ubuntu installation files. This isn't very complicated to do, but may not work very well with driver files because of the subtle differences in where system files are stored across different Linux distributions. You can find more information about `alien` at www.kitenet.net/programs/alien/. It's contained within the Ubuntu software repositories and can be downloaded using the Synaptic Package Manager, as explained in the "Setting Up Online Software Repositories" section later in this chapter.

Proprietary vs. Open Source Drivers

As discussed earlier in this book, Linux is an open-source operating system. This means that the source code underlying Linux programs is available for inspection. This is a good thing when it comes to hardware drivers, because bugs in the code can be spotted and repaired by anyone with an interest in doing so. If you consider that a bug in a graphics driver could mean your PC crashes every five minutes, the value of such an approach is abundantly clear.

Unfortunately, some hardware manufacturers don't like to disclose how their hardware works, because they want to protect their trade secrets. This makes it impossible for them to release open-source drivers, because such drivers would expose exactly how the hardware operates. Such companies are not blind to the fact that growing numbers of people use Linux, so they release proprietary drivers, whose source code is not made publicly available.

Note Sometimes, it's actually possible to use Windows drivers under Linux, such as with the `NdisWrapper` project (<http://ndiswrapper.sourceforge.net/>).

Proprietary drivers bring with them a number of problems. The first relates to bugs. To use a proprietary driver is to be at the mercy of the hardware manufacturer's own development and release schedule. If the driver has a serious bug, you'll either have to work around it or put up with troubling issues until the manufacturer offers an update. Additionally, proprietary

drivers are usually tied into particular versions of Linux. Sometimes, they can be hacked to work with other versions of Linux, but this isn't something that beginners can take on themselves.

Although the folks behind Ubuntu strongly support open-source software, even they realize proprietary drivers need to be used in certain situations. For example, it's impossible to use the 3D graphics elements of most modern graphics cards unless you have a proprietary driver. Because of this, it's often possible to grab proprietary drivers from the Ubuntu online software repositories. We'll look at connecting to these software repositories later in this chapter, in the "Setting Up Online Software Repositories" section, and you'll also learn how to download 3D graphics card drivers, in the "Installing 3D Graphics Card Drivers" section.

Note Linux sees hardware in a technical way, rather than in the way humans do. If you attach something like a USB CD-R/RW drive, Linux will recognize the drive hardware and attempt to make it work. It won't try to find a driver for that specific make and model of CD-R/RW drive. Thus, Linux is able to work with a wide range of hardware because a lot of hardware is actually very similar on a technical level, despite the differences in case design, model names, and even prices!

WHAT HARDWARE WORKS?

The question of what hardware works under Ubuntu is one that's not easily answered. However, you can take a look at <http://doc.gwos.org/index.php/HCL> to see if your hardware is listed. This is an informal list created by the people who run Ubuntu's online forums, and it's not comprehensive (which is to say that there may be hardware that works fine that isn't mentioned). Nor is the list guaranteed to be 100% accurate. But it's certainly worth a look.

You can find a Ubuntu hardware database at <http://hwdb.ubuntu.com>, but at the time of writing, it's fairly primitive and shows only basic details about various users' systems. However, you can help the effort by submitting your own system information to the list: click Applications ► System Tools ► Ubuntu Device Database. This will anonymously submit a list of your computer's hardware. It's very likely that the list will be opened up to the public very soon, so you should certainly check the site if you're thinking of adding any new hardware to your system.

A search engine like Google is your best friend if the two Ubuntu hardware lists don't help. Simply search for the brand and model of your hardware and add "Ubuntu" to the search string. This should return results, usually from the Ubuntu forums (www.ubuntuforums.org) or an individual's blog, written by those who have found a way to make that type of hardware work.

Viewing Your Hardware

When using Windows, you might have come across Device Manager, the handy tool that lists your PC's hardware. Ubuntu contains a similar piece of software, as shown in Figure 8-1, which you can open by selecting System ► Administration ► Device Manager.

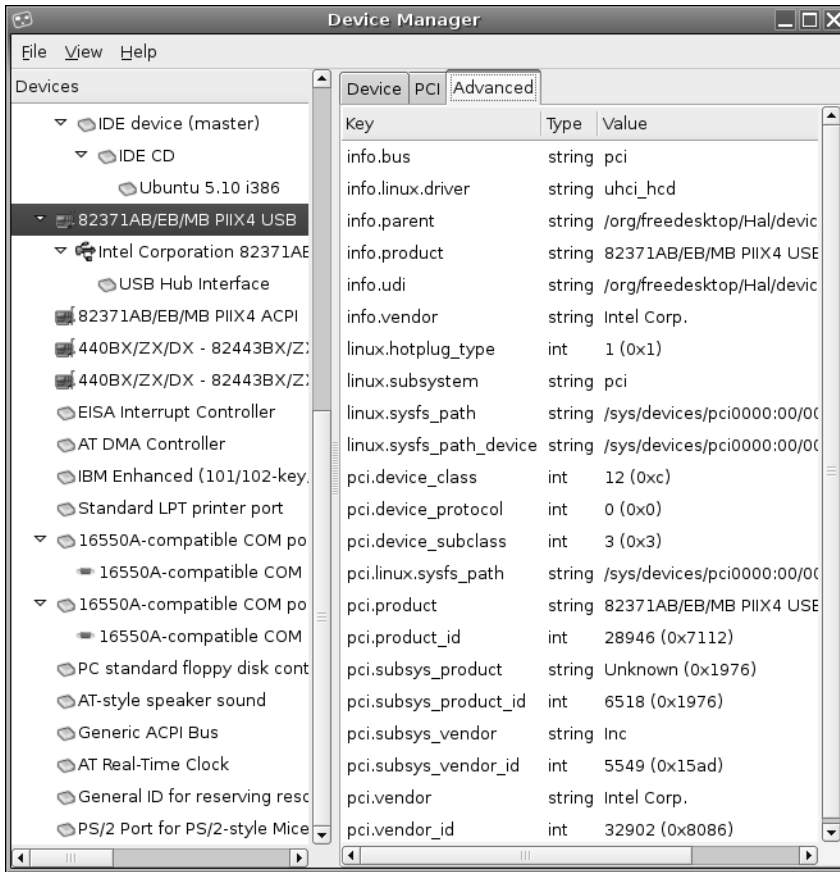


Figure 8-1. Ubuntu's Device Manager program can display just about everything you need to know about attached hardware.

You should be aware of a few important differences between the Windows and Ubuntu versions of Device Manager. Under Ubuntu, the list is for information only. You can't tweak any settings and must instead rely on separate pieces of configuration software. On the other hand, Ubuntu's list is far more comprehensive than that in Windows. In Ubuntu, Device Manager thoroughly probes the hardware to discover its capabilities.

Perhaps the biggest difference, however, is that just because a piece of hardware is listed within Ubuntu's Device Manager doesn't mean that the hardware is configured to work with Ubuntu. In fact, it doesn't even imply that the hardware will *ever* work under Ubuntu. Device Manager's list is simply the result of probing devices attached to the various system buses (PCI, AGP, USB, and so on) and reporting back the data.

Nonetheless, Device Manager is the best starting place if you find that a certain piece of hardware isn't working. If a piece of hardware is listed, then it proves, if nothing else, that the system recognizes that the hardware is attached.

That said, it's very unlikely that you'll need to use Ubuntu's Device Manager, because most hardware will work after you install Ubuntu. At the very most, all you'll need to do is configure a few settings for the hardware.

Note To configure most hardware devices, you'll be asked for your password, because configuring hardware requires superuser privileges.

Getting Online

Getting online is vital in our modern Internet age, and Ubuntu caters to all the standard ways of doing so. Linux was built from the ground up to be an online operating system and is based on Unix, which pioneered the concept of networking computers together to share data back in the 1970s. However, none of this is to say that getting online with Ubuntu is difficult! In fact, it's very easy.

Regardless of whether you use a modem, standard Ethernet card, or wireless card, the same program is used to configure your network settings under Ubuntu. Support for many makes and models of equipment is built in, so in most cases, all you need to do is enter a few configuration details.

Note Linux actually runs around 60% of the computers that make the Internet work! Whenever you visit a web site, there's a strong chance that it's run using Linux. As your Linux skills increase, you'll eventually get to a stage where you, too, can run your own Internet servers. It sounds difficult, but can be quite easy.

Using an Ethernet Card

Ethernet is one of the oldest and most established network technologies. When we talk of Ethernet, we refer to wired networks—all the computers on the network are connected by cabling to a central hub or router. (The other form of networking technology, which works without wires, is covered in the next section.)

You might go online via an Ethernet card in a variety of situations. If you have DSL or cable broadband service at your home or workplace, for example, you might use a DSL router. Your computer will then connect to this router via Ethernet, and all you need to worry about on your PC is getting your Ethernet card up and running.

Note Using a DSL or cable modem router is the preferred way of going online via broadband. However, some people use USB-based DSL modems, which connect to and are operated by their PC. These are covered in the “Using a USB-Based DSL Modem” section later in this chapter.

If you're running Ubuntu on a PC in an office environment, it's likely that you will connect to the local area network using an Ethernet card. This lets your computer communicate with other computers, as well as with printers. In some offices in which an Internet connection is provided, this connection will allow you to go online.

Configuring a Network Card via DHCP

Most computers that connect to a broadband router or an office network receive their configuration data via the Dynamic Host Control Protocol (DHCP), which is to say that your computer receives its IP, gateway, subnet mask, and DNS addresses automatically. However, you will need to configure your network card to work via DHCP. You can do this configuration with the Network Settings applet, which can be found under the System ► Administration menu. Because you're going to configure hardware settings, you'll need to enter your password to proceed.

Follow these steps to configure your network card:

1. Select System ► Administration ► Networking to open the Network Settings dialog box.
2. You should find your Ethernet card at the top of the list. It will be identified as eth, followed by a number, such as eth0 or eth1. Ubuntu should indicate that the card isn't configured. Click the entry for your Ethernet card, and then click Properties.

Note Listed beneath the Ethernet card will be any other networking devices you might have, such as a dial-up modem. If you don't want to use these, you can leave them unconfigured.

3. In the Properties dialog box, put a check in the Enable This Connection check box, and make sure the Configuration drop-down list reads DHCP. Then click OK.
4. In the Network Settings dialog box, click Activate. Then click OK.
5. After a few seconds, your network card should be up and running. Test it by using the web browser to visit a site.

From this point on, your network card should automatically activate each time you boot, so you should not need to return to the Network Settings applet.

Configuring a Static IP Address

On some networks, you might have been assigned an IP address, which you must enter manually, along with a few other networking addresses. This is referred to as a *static IP address*.

You should speak to your system administrator or technical support person to determine these settings. Ask the administrator for your *IP address*, *DNS server addresses* (there are usually two or three of these), your *subnet mask*, and the *router address* (sometimes referred to as the *gateway address*). The settings you will get from your system administrator will usually be in the form of a series of four numbers separated by dots, something like 192.168.0.233.

Once you know your settings, proceed as follows.

1. Select System ► Administration ► Networking to open the Network Settings dialog box.
2. Find your network card in the list (it should be referred to as eth0), click its entry, and then click Properties.

3. In the dialog box that appears, put a check in the Enable This Connection check box. In the Configuration drop-down list, make sure Static IP Address is highlighted. In the IP Address, Subnet Mask, and Gateway Address text boxes, fill in the relevant details. Figure 8-2 shows an example of a completed Interface Properties dialog box. Click OK after filling in the information.

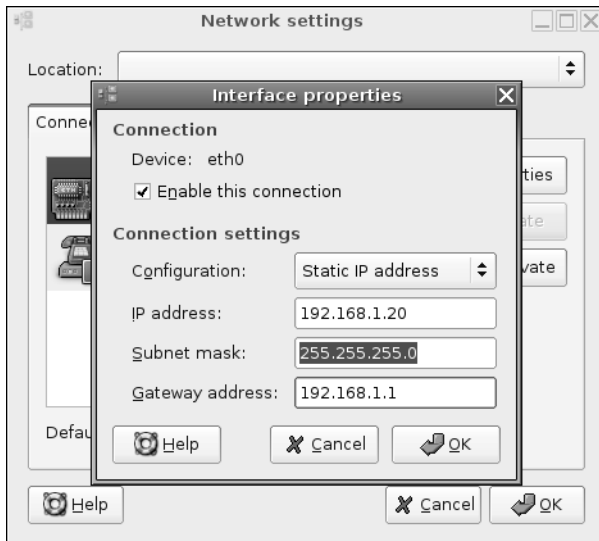


Figure 8-2. You can configure Ubuntu to work with DHCP networks, or you can define a static IP address.

4. In the Network Settings dialog box, click the DNS tab.
5. Click the Add button, and then type the first DNS address. Click Add again, and enter the second DNS address, if you have been given one (and then enter a third if you were given one). Click OK.

Tip If you're using a static IP address with a router, such as that provided by a DSL modem, the DNS address will probably be the same as the router/gateway address.

6. Click the Connections tab in the Network Settings dialog box.
7. Click the Apply button, and then click OK.

Your network connection should now work. Test it by using the web browser to visit a web site. If you find it isn't working, try rebooting. However, if your system administrator mentioned that a proxy must also be configured, you'll also need to follow the instructions in the "Working with a Proxy Server" section later in this chapter.

Joining a Wireless Network

A wireless (also referred to as Wi-Fi) network is, as its name suggests, a network that does away with cabling and uses radio frequencies to communicate. It's more common for notebooks and handheld computers to use wireless connections, but some desktop computers also do.

Notebooks and PDAs typically use built-in wireless network cards, with an antenna built in to the case. However, some notebooks might use PCMCIA cards, which will have an external square antenna, as shown in Figure 8-3.



Figure 8-3. *If you use a PCMCIA wireless card, it will most likely have an external antenna.*

Configuring a Wireless Card

The process for configuring a wireless network card is quite similar to configuring a standard Ethernet card, meaning you can follow the instructions for configuring a network card in the previous section, with the following differences:

- The wireless card will be identified as *Wireless Connection*. If your computer also has a standard Ethernet adapter in addition to wireless capabilities, you'll need to make sure you select the wireless card when choosing an adapter to configure.
- You need to select the base station you wish to connect to in the Network Name (ESSID) drop-down list or, if your base station is configured so it doesn't broadcast its name, you should type it manually into the Network Name (ESSID) field. If you work in an office environment, you may need to speak to your network administrator to find out the base station name.

- If the connection doesn't use Wireless Equivalent Privacy (WEP) protection, you can leave the WEP Key field empty (it doesn't matter what the Key Type drop-down list reads). If it does use WEP protection, continue reading for instructions on how to configure it.

In most instances, wireless network cards are configured with DHCP, so that they grab a network address automatically. The nature of a wireless network, where many people might join or leave the network at will, means that using static IP addresses is a bad idea.

Caution When I configured a wireless card on a desktop computer, I found that I was unable to use DHCP. The card would pause for a long time without being able to find an address. The fix was to assign a static IP address to the card. I simply examined my router's settings to find the range of IP addresses that it normally assigns via DHCP and chose one to give my PC. This occasional inability to get an address by DHCP is clearly a bug that may have been fixed by the time you read this, so consider performing a system update as soon as you're online (see Chapter 9).

Using WEP or WPA Encryption

Some wireless networks use the WEP or Wi-Fi Protected Access (WPA) system. These systems encrypt the data being transmitted on the network so it cannot be stolen by hackers with special equipment. Also, people can't join the wireless network unless they know the encryption key, which is basically an access code. This prevents unauthorized people from accessing the network.

Note At one time just a few years ago, people walking or driving past office blocks would be able to steal a business's Internet connection! This practice, known as *war driving*, became a hobby for some people. It's still possible today, although businesses have realized the dangers and usually protect their networks. Additionally, connecting to a wireless network connection without permission has been made illegal in some countries.

The encryption key normally takes the form of a string of letters and numbers, which you should get from your system administrator. Alternatively, your administrator may give you a passphrase, which might be a sentence in English, including spaces between words. You then enter that as a kind of long password.

As of this writing, WEP encryption is supported in Ubuntu, while WPA support is still in its infancy. It's possible to get WPA working with a little hard work, but it's not something for beginners.

WEP is by far the easiest choice when it comes to configuration. Although WEP has been found to be relatively easy to circumvent, for most home users, it offers sufficient protection. If possible, you should reconfigure your wireless base station to use WEP rather than WPA, and configure Ubuntu accordingly. Only if you have no choice, or if you're sufficiently confident of your Linux command-line skills, should you configure WPA.

SECURE SOCKET LAYER (SSL)

The transfer of confidential or financial data across the Web—to and from online banking sites, for example—is nearly always protected by an additional layer of secure protection called Secure Socket Layer (SSL) HTTP. You can tell if this is the case with any site because the address will begin with `https://`. Additionally, most browsers display a padlock symbol at the bottom of the screen. Accessing such sites should be safe, even if your wireless connection isn't protected with either WEP or WPA.

Similarly, although online shopping sites don't use SSL while you're browsing, when it's time to pay, they always use SSL. This ensures your credit card details are encrypted. If the store doesn't adopt an `https://` address when you click to visit the virtual checkout, you shouldn't shop there!

So, there's an argument to be made that you don't really need WEP or WPA protection if you simply use your wireless connection to browse the Internet. However, some web mail services transfer your username and password "in the clear," which is to say without using SSL. This means this information could be picked up by an eavesdropper. In the case of Hotmail and Yahoo Mail, you can select secure login, but it isn't activated by default. Google Mail appears to use SSL all the time for login, but after this, your e-mail messages are transmitted across the Internet in the clear and, in theory, can be eavesdropped by anyone, anywhere.

Configuring WEP

WEP keys come in either hexadecimal (hex) or plain text (passphrase) varieties. Hex keys look like this in their 128-bit form: CB4C4189B1861E19BC9A9BDA59. In their 64-bit form, they will be shorter and may look similar to 4D9ED51E23. A passphrase will take the form of a single short sentence. Ubuntu can work with both 64- and 128-bit keys, as well as passphrases.

Follow these steps to configure WEP (these steps assume your network card has already been configured, as described above):

1. Select **System** ► **Administration** ► **Networking** to open the Network Settings dialog box.
2. Double-click the entry in the list corresponding to your wireless card.
3. In the Key Type drop-down list, select Hexadecimal, if you have a hex key, or select Plain (ASCII), if you have a passphrase. In the WEP Key box, type the key or passphrase. Then click OK.
4. In the Network Settings dialog box, click the entry in the list for your wireless card, as shown in Figure 8-4. Select Deactivate, and then click Activate. Then click OK.
5. Test your connection by attempting to browse to a web site.

If you find your card doesn't seem to work after you enter new WEP settings, try rebooting. Then open the Network Settings applet and make sure your card is activated.

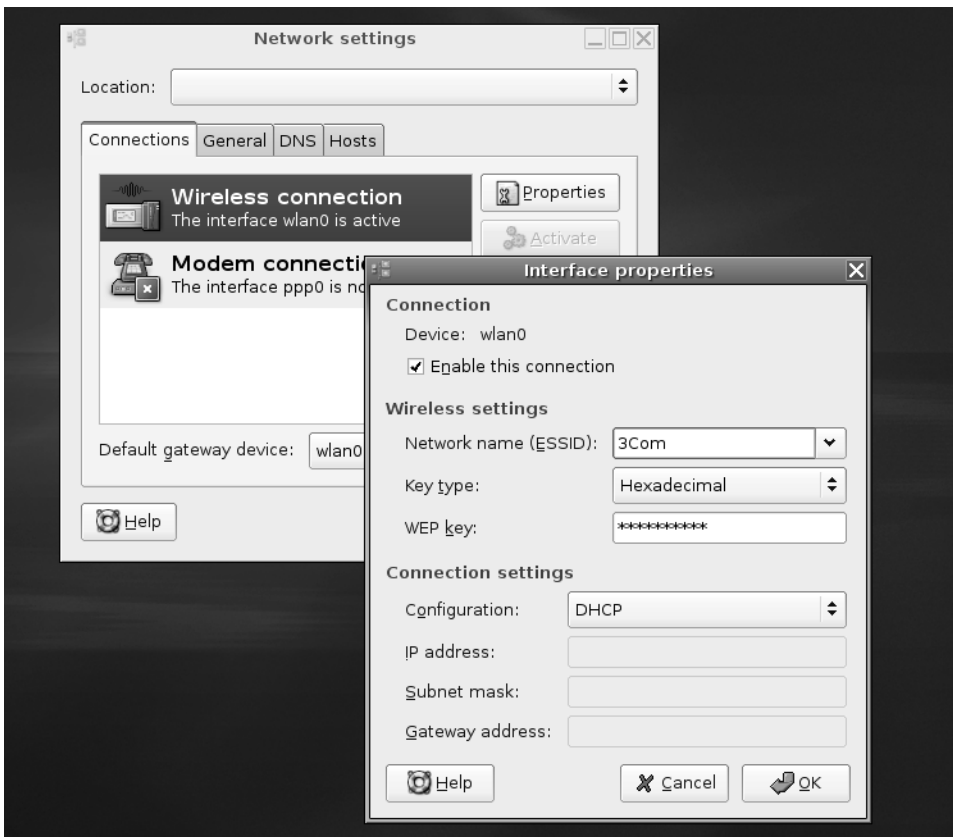


Figure 8-4. Ubuntu is able to join WEP-protected wireless networks, using either 64- or 128-bit keys.

Configuring WPA

As I noted earlier, support for WPA has yet to be officially added to Ubuntu, but it's possible to join networks using WPA—with a little work. This involves downloading additional software from the Ubuntu software repository and editing a couple of configuration files. This configuration should be undertaken only by experienced Ubuntu users, since it involves heavy use of the command-line prompt.

Note It's very likely that WPA will be officially supported by the Ubuntu project before long and, in fact, perhaps by the time you're reading this book. If it is supported, it should become an option within the Network Settings applet alongside WEP. Adding support for new technologies like WPA is one of the reasons you will want to keep your system up-to-date (see Chapter 9).

Perhaps it goes without saying that you should ensure your wireless card is compatible with WPA before you start. It's found on models of one or, at most, two years old. You can learn this information from the card's packaging or by visiting the manufacturer's web site.

Caution Some earlier wireless cards rushed to include WPA before it was fully certified and may feature buggy implementations. In such cases, it might be possible to update the firmware of the card, usually using a utility that runs in Windows. If you suspect your card has this fault, check the manufacturer's support site for details.

You'll also need to know if the card uses a driver module that's compatible with `wpa_supplicant`, the software that handles WPA management under Ubuntu. As this book went to press, `wpa_supplicant` was compatible with the following wireless drivers: HostAP, Prism54, Madwifi, Atmel, `wext`, `NdisWrapper`, and `ipw`.

You can discover which driver your card uses through Ubuntu's Device Manager. Open Device Manager (System ► Administration ► Device Manager) and look for your wireless card in the list of hardware on the left side of the window. Once you find it, select it, and then click the Advanced tab on the right side of the window. Find the line that reads `info.linux.driver`, and look in the Value column for the name of the driver. If what you discover isn't in the list of compatible drivers, then `wpa_supplicant` won't work with your card. If it is present, make a note of the name of the driver, because you'll need to know it to complete the configuration.

Finally, you'll need to know how Ubuntu refers to your wireless card. Open the Network Settings applet (System ► Administration ► Networking) and look for the entry in the list referring to your wireless card. It should read, "The interface for *x* is active/not configured," where *x* will be `eth0`, `ath0`, or `wlan0` (or possibly something else). Whatever the case, make a note of it for later use.

Note There's a chance that `wpa_supplicant` will support more drivers than those I've listed by the time you read this. With the version I used, I could find out which drivers were supported by viewing the help documentation that appeared when I typed `wpa_supplicant`, without any command options, in a GNOME Terminal window.

To install the `wpa_supplicant` software, first configure the Synaptic Package Manager to access Ubuntu's online software repositories. This is necessary to download the `wpa_supplicant` software package. If you haven't already set up online repositories, see the "Setting Up Online Software Repositories" section later in this chapter. Once that's completed, return here.

Obviously, you'll need to find some way of getting online to do be able to download using the Synaptic Package Manager. Consider making a direct Ethernet connection to your router, or simply turn off WPA on the router for the duration of the download.

Follow these steps to configure WPA:

1. Select System ► Administration ► Synaptic Package Manager.
2. Click the Search button. Enter `wpa_supplicant` as a search term.

3. In the list of results, click the check box next to `wpa_supplicant` and `opt` to install it.
4. Open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type the following to open in Gedit the first `wpa_supplicant` configuration file you'll need to edit:

```
sudo gedit /etc/default/wpa_supplicant
```

5. Look for the line that reads `ENABLED=0`. Change it to `ENABLED=1`.
6. Beneath this, you'll see a line that reads `OPTIONS="-w"`. Change this to read as follows:

```
OPTIONS="-D drivervname -i adaptername -c /etc/wpa_supplicant.conf -w"
```

Replace *drivervname* with the name of the driver you discovered earlier. Enter it in lowercase (for example, `prism54` or `atmel`). Replace *adaptername* with the name Ubuntu assigns your card, which you also discovered earlier (such as `eth0`, `ath0`, `wlan0`, or something similar).

7. Save the file and close Gedit.
8. Now you need to edit the second `wpa_supplicant` configuration file. In the GNOME Terminal window, type the following:

```
sudo gedit /etc/wpa_supplicant.conf
```

9. At the bottom of the file will be four lines, the first one of which will read `network={` and the last of which will be a `}` symbol by itself. At the beginning of each of these lines, insert a `#` symbol. In other words, the first line will now read `#network={`, and the last will read `#}`. Then save the file.
10. It's now time to enter the WPA passphrase. Typing the following will immediately prompt you for the passphrase and also automatically update your configuration files:

```
wpa_passphrase wifi_station_name | sudo tee -a /etc/wpa_supplicant.conf
```

Replace *wifi_station_name* with the SSID name of your wireless base station. The SSID name is how the base station identifies itself and, provided the base station is set to broadcast its name, you can discover it by using the Network Settings applet and looking in the Network Name (ESSID) drop-down list for your wireless card.

11. The cursor will move to a blank line. Type the WPA passphrase for your router, and then press Enter.
12. Now you need to make `wpa_supplicant` start at bootup. Type the following in the terminal window:

```
sudo ln -s /etc/init.d/wpa_supplicant /etc/rc2.d/S40wpa
```

13. Reboot, and you should now find that you are able to connect to the WPA-protected wireless network.

USING NDISWRAPPER

You might find that your wireless card is not supported under Ubuntu. While the people who write drivers for Linux endeavor to have the operating system work with all wireless cards, they often need to reverse-engineer wireless card technology, and this can mean that drivers take a while to appear. In cases such as this, it's sometimes possible to use NdisWrapper. This allows you to use Windows XP drivers for the wireless card. It's not a perfect solution, but it's certainly worth a try.

Sadly, although NdisWrapper is fundamentally simple to use, there are special considerations for each card and it's not possible to provide an all-encompassing guide here. The best policy if you wish to use NdisWrapper is to search the Ubuntu Forums (www.ubuntuforums.org) for a guide to configuring your particular card. It's almost certain that at least one other person will have posted a how-to tutorial. However, you should first check <http://ndiswrapper.sourceforge.net/mediawiki/index.php/List> to see if your card is listed among those known to work.

Configuring Dial-Up Internet Access

At one time, nearly all home users (and even some workplaces) used dial-up modems across phone lines to access the Internet. But since the boom in broadband courtesy of DSL and cable television, dial-up access has declined in popularity.

Ubuntu has excellent support for dial-up connections but, unfortunately, only for certain modems. In simple terms, it supports many external serial modems and some internal models (including PCI and PCMCIA cards), but it doesn't directly support so-called Winmodems, which are the most popular type of dial-up modem on the market today.

Winmodems rely on Windows taking care of the coding/decoding work, hence the name. This is done to save on component costs during manufacturing.

The companies that make the hardware behind Winmodems aren't always helpful to the Linux kernel programmers, and it can be difficult for them to make the modems work under Linux. Yet more manufacturers provide the required drivers but under a software license that's not compatible with the GNU Public License (GPL) used throughout the rest of Linux. This means that including the software with Linux brings with it ethical and even legal issues.

A project that works to bring support for most Winmodems to Linux is called Linmodems. However, the Ubuntu developers have decided not to include its software within Ubuntu at the moment. If you would like to try to make your Winmodem work under Ubuntu, visit the Linmodem project's web page, <http://linmodems.org>. You'll find a welcoming community of users who are dedicated to helping others get online. Alternatively, you can skip straight to the guide that's written for beginners who would like to make a Winmodem work under Linux: <http://linmodems.technion.ac.il/first.html>. Searching the Ubuntu forums (www.ubuntuforums.org) might also reveal useful information.

Note The easiest solution if you have a Winmodem and intend to use dial-up connections on your computer is to buy an external modem that connects to your computer via its serial port. These types of modems can be purchased for as little as \$10 to \$20, if you shop around. Auction sites like eBay (www.ebay.com) are a particularly good place to look for external modems. Practically all external modems are guaranteed to work with Ubuntu.

Before you start configuring your dial-up modem, you'll need to know your dial-up username and password, and also the phone number to use in order to get online. You should ensure that your modem is plugged into the phone socket and is powered up.

1. Right-click a blank spot on the Panel at the top of the Ubuntu desktop and select Add to Panel.
2. In the list of applets, scroll down to the System & Hardware section and click Modem Monitor. Then click the Add button.
3. You should find a new phone icon on your Panel. Right-click it and select Properties.
4. In the window that appears, put a check in the Enable This Connection box, and then fill in the details in the rest of the dialog box.
5. In the Phone Number field, type the ISP's dial-up number, including any area code or other prefix.
6. In the Dial Prefix field, type the number that you might need to access an outside line if you're in an office. Alternatively, you might choose to enter the prefix you normally dial to activate a particular call plan. Otherwise, this box can be left empty.
7. In the Username and Password fields, enter your dial-up username and password (*not* your Ubuntu login username and password!).
8. Click the Modem tab and, if you have an external modem, select `/dev/ttyS0` in the Modem Port drop-down list. This assumes that the modem is plugged into the first of your serial ports. If you know that it's plugged into the second serial port, select `/dev/ttyS1`. If your modem is internal, you may need to select `/dev/modem` or `/dev/ttyS4`, depending on your hardware configuration. If you find your modem connection doesn't work when you test it, come back to this step and try a different port setting.
9. If you wish to hear your modem dialing and making the connection, click the Volume drop-down list.
10. Click the OK button.

After this setup, you can connect to your ISP by right-clicking the phone icon and selecting Activate, as shown in Figure 8-5. When you wish to hang up, right-click and select Deactivate. Note that you'll be prompted to enter your Ubuntu login password whenever you attempt to go online.

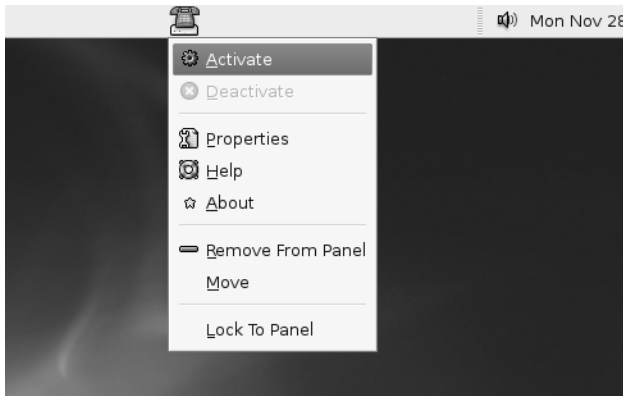


Figure 8-5. Going online with your dial-up modem is simply a matter of right-clicking the icon and selecting *Activate*.

Using a USB-Based DSL Modem

Some people connect to broadband Internet services using a DSL modem that connects to the computer via USB. This is very common in the United Kingdom, where such modems are provided free of charge as part of ADSL broadband subscriptions.

Sadly, as with dial-up modems, support for the majority of USB-based DSL modems isn't included in Ubuntu. To get such modems working, you must download and install additional drivers, as well as configuration software. This can be rather complicated and, if you can afford it, by far the best policy is to upgrade to a router-based modem and connect via an Ethernet or wireless connection. Such modem routers are inexpensive nowadays, and most modern PCs have Ethernet ports built in. The popular HowStuffWorks site (<http://computer.howstuffworks.com/home-network.htm>) produces a good guide for beginners that explains home networking.

Here, I'll explain how to configure the BT Voyager series, which is common in the United Kingdom, and the Alcatel SpeedTouch series, shown in Figure 8-6, which is more common in the United States and Europe. The BT Voyager uses the Globespan chipset, and with some minimal adaptation, the instructions provided here should work for all modems based on this hardware. See <http://eciadsl.flashtux.org/modems.php?lang=en> for a complete list. On this web page, you'll also find links to software that supports other USB ADSL modems, such as those made by Olitec, Alcatel, and Elsa (follow the Conexant link for Olitec).

Prior to following these steps, you'll need to find out your DSL technical settings, including the VPI/VCI settings as well as your username and password. These can be discovered by phoning your ISP's technical support line or by looking through the documentation sent to you when you subscribed to the service.



Figure 8-6. If you use an Alcatel SpeedTouch USB modem, you can get online with just a little configuration work.

Setting Up a BT Voyager USB ADSL Modem

The instructions in this section will work with all USB models of the Voyager modem, up to Voyager 105. If you have a Voyager 205, which includes both USB and router functionality, you're strongly advised to connect via Ethernet instead of USB. This simply involves running a Cat 5 cable from your PC's Ethernet port to one of the Ethernet inputs on the Voyager modem, and then following the instructions to configure a DHCP Ethernet connection, presented earlier in this chapter.

You'll need to download a package file, so you must go online by some other means, perhaps via dial-up (see the previous section). Alternatively, you might grab the files on a different computer and transfer them onto your Ubuntu system via a floppy disk or USB memory stick (see the "Using a USB Memory Stick" section later in this chapter).

Follow these steps to set up a BT Voyager modem:

1. Visit <http://eciadsl.flashtux.org/download/> and download the `eciadsl-synch_bin.tar.bz2` file to your `/home` directory.
2. At the same download site, grab the latest version of the `eciadsl-usermode` file that has an `.i386.deb` file extension. The version numbers appear in the middle of the filename. At the time of writing, the latest version was `eciadsl-usermode_0.11-1.i386.deb`. Create a new directory within your `/home` directory called `packages` and download it there.
3. Visit <http://archive.ubuntu.com/ubuntu/pool/universe/r/rp-pppoe/> and download the latest version of the `pppoe` file (*not* `rp-pppoe!`) that has an `.i386.deb` file extension. At the time of writing, this was `pppoe_3.5-4ubuntu1_i386.deb`. Save it into the `packages` directory you created earlier.

4. Open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type the following in sequence (the ~ symbol can be found above the # key on UK keyboards and above the top-left key on US keyboards). Note that you'll need to type your login password when prompted.

```
cd ~/packages/
sudo dpkg -i *
cd ~
tar xjf eciadsl-synch_bin.tar.bz2
sudo mv eciadsl-synch_bin/* /etc/eciadsl/
rm -rf eciadsl-synch_bin/
```

5. To run the modem configuration program, type the following:

```
sudo eciadsl-config-tk
```

6. You'll see the window shown in Figure 8-7. Start by entering your username and password (if required) in the User and Password fields.




Figure 8-7. It's possible to get USB-based DSL modems working with a little effort.

7. Enter the VPI and VCI values in the fields at the top right of the program window.
8. In the Select Your Modem list, click the entry relating to your modem.

9. Put a check alongside Update Provider DNS and select your ISP from the list.
10. In most cases, the Select a PPP Mode, Chipset, and Change Synch .bin File settings can be left unchanged. However, if you use a Voyager 105 modem, put a check in the box alongside Change Synch .bin File, and then select `gs7470_synch03.bin` from the list.
11. Click the Create Config button.
12. Back in the GNOME Terminal window, type the following:

```
sudo eciadsl-start
```
13. This will connect you to your ISP via the modem. Attempt to browse a web site to see if the connection works.
14. In order for your connection to start automatically every time you boot, you will need to create a startup script. In the GNOME Terminal window, type the following:

```
sudo gedit /etc/init.d/voyager_start
```
15. This will open a Gedit window in which you should type the following, and then press Enter:

```
eciadsl-start
```
16. Save the file.
17. In the GNOME Terminal window, type the following:

```
sudo chmod +x /etc/init.d/voyager_start
sudo ln -s /etc/init.d/voyager_start /etc/rc2.d/S20voyager_start
```
18. Finally, delete the package directory you created earlier, as well as the `eciadsl-synch_bin.tar.bz2` file.

Don't forget that logging on via DSL can take some time, depending on the quality of your connection. On my test system, I found that it could take up to a minute after the desktop appeared to establish a connection. I also found the modem wouldn't work unless other network interfaces, such as Ethernet, were deactivated via the Network Settings applet (it was necessary to double-click the network card's entry in the list and remove the check from Enable This Connection). I also found that I had to unplug and reattach the modem before undertaking step 12 in order to reset the modem.

Setting Up an Alcatel SpeedTouch Series USB DSL Modem

Before commencing configuration of an Alcatel SpeedTouch modem, you'll need to find out if your ISP uses the PPPoE or PPPoA technical standard. You can discover this by browsing to <http://linux-usb.sourceforge.net/SpeedTouch/faq/index.html#q12> and searching for your country's entry in the table. If your country is listed as using both standards, you should follow the instructions for PPPoA in the following steps.

You'll need to download several files to complete this configuration, so you will either need to get your computer online by some other method, such as dial-up, or grab the files mentioned on a different computer and transfer them across to Ubuntu via floppy disk or

USB memory stick (see the “Using a USB Memory Stick” section later in this chapter). The SpeedTouch driver has two components, and you’ll need to download both.

Follow these steps to configure your modem (these instructions are adapted from a guide that you can find at <http://linux-usb.sourceforge.net/SpeedTouch/ubuntu/warty.html>):

1. Download the first driver file by browsing to www.speedtouch.com/download/drivers/USB/SpeedTouch330_firmware_3012.zip in your browser. Place the downloaded file in your `/home` directory.
2. Visit http://linux-usb.sourceforge.net/SpeedTouch/ubuntu/modem_run. This will download the GPL component of the driver. Again, place the downloaded file in your `/home` directory.
3. If your ISP uses PPPoE, you’ll also need to visit http://linux-usb.sourceforge.net/SpeedTouch/ubuntu/br2684ctl_20040226-1_i386.deb and download an additional file. Place it in your `/home` directory.
4. Because different drivers are used for each revision of the hardware, you will need to discover which version of the modem firmware you have. Open a GNOME Terminal window (Applications ► Accessories ► Terminal), and then type the following two lines, hitting the Enter key after each (the straight line is called a *pipe symbol*, which you can type by pressing Shift+\):

```
cat /proc/bus/usb/devices | grep -B 1 THOMSON
cat /proc/bus/usb/devices | grep -B 1 ALCATEL
```

In the results will be a line that reads, at the end, Rev=x, where x is the firmware version number. Make a note of it for later.

5. Make sure the files you downloaded earlier are in your `/home` directory, and then type the following in the GNOME Terminal window to copy the firmware files to your system directories (the `~` symbol can be found above the # key on UK keyboards and above the top-left key on US keyboards). Note that you’ll need to type your login password when prompted.

```
cd ~
unzip SpeedTouch330_firmware_3012.zip
```

If your modem uses firmware version 4, type the following

```
sudo cp ZZL_3.012 /etc/ppp/mgmt.o
```

Otherwise, type the following:

```
sudo cp KQD6_3.012 /etc/ppp/mgmt.o
```

6. Now you must copy the `modem_run` file you downloaded earlier to your system directories and make it executable:

```
cd ~
sudo cp modem_run /usr/sbin
sudo chmod +x /usr/sbin/modem_run
```

7. If your ISP uses PPPoE, you'll need to install the `br2684ctl_20040226-1_i386.deb` file as well:

```
sudo dpkg -i br2684ctl_20040226-1_i386.deb
```

8. Now you must add your login details to the files `pap-secrets` and `chap-secrets`. The following will open both files in Gedit on two separate tabs:

```
sudo gedit /etc/ppp/pap-secrets /etc/ppp/chap-secrets
```

9. In *both files*, add the following line to the bottom:

```
"username" * "password"
```

Replace *username* and *password* with your ISP login details. For example, my username for my ISP is `keir.thomas@btbroadband.com`. My ISP doesn't use passwords, so I simply typed "" in the password area. If you find this doesn't work when you test this connection, try repeating this step and typing any word between the password quotation marks; it doesn't matter what you type.

10. Save each file, and then exit Gedit.
11. Type the following to create the `speedtch` configuration file and edit it with Gedit:

```
sudo gedit /etc/ppp/peers/speedtch
```

12. Add the following lines:

```
noipdefault
defaultroute
user 'username'
noauth
updetach
usepeerdns
plugin rp-pppoe.so
nas0
VPI.VCI
```

Replace *username* with your ISP username, and replace *VPI.VCI* with your VPI and VCI settings, separated by a period. For example, because I live in the United Kingdom, I typed `0.38`. Your ISP's technical support line will be able to supply these values. (Note that VPI/VCI values are sometimes referred to as *VP/VC*.)

13. Save the file.
14. Type the following at the GNOME Terminal command-line prompt:

```
sudo rm -f /etc/resolv.conf
sudo ln -s /etc/ppp/resolv.conf /etc/resolv.conf
```

15. Next, you need to ensure certain modules load during boot. Load the `modules` file into Gedit:

```
sudo gedit /etc/modules
```

- 16.** Add the following lines to the bottom of the file:

```
ppp_generic
pppoatm
```

If you're using PPPoE, you'll also need to add `br2684` to the end of the file.

- 17.** Now all that remains is to configure the modem to connect to your ISP upon bootup. To do this, you'll need to create a startup script. Type the following to create the file and edit it in Gedit:

```
sudo gedit /etc/init.d/alcatel_connect
```

If you're using PPPoA, type the following into the file, pressing Enter after the last line:

```
modem_run -k -f /etc/ppp/mgmt.o &&
pppd call speedtch
```

If your ISP uses PPPoE, type the following:

```
modem_run -k -f /etc/ppp/mgmt.o &&
sleep 10 &&
br2684ctl -b -c 0 -a VPI.VCI &&
sleep 20 &&
pppd call speedtch
```

Replace *VPI.VCI* with your VPI and VCI values, separated by a period. For example, I typed `0.38`.

- 18.** Save the file.
- 19.** Type the following in the GNOME Terminal window:

```
sudo chmod +x /etc/init.d/alcatel_connect
sudo ln -s /etc/init.d/alcatel_connect /etc/rc2.d/S95alcatel_connect
```

- 20.** Reboot, and you should find that your modem automatically connects and goes online.

Logging on via DSL can take some time, depending on the quality of your connection. You might find that a connection won't be established for up to a minute after the desktop appears. The LED lights on your modem should indicate the state of the connection.

You may find that the modem won't work unless other network interfaces, such as Ethernet, are deactivated via the Network Settings applet (if necessary, double-click the network card's entry in the list and remove the check from the Enable This Connection check box).

Working with a Proxy Server

Some networks in offices require that you use a web proxy (often referred to as an *HTTP proxy*). A proxy is a server computer that does two things. First, it provides additional security by providing a single portal to all web pages. Second, it helps speed up Internet access by storing frequently accessed pages. This means that if ten people request the same web page, there's no need to get the same ten pieces of data from the Internet. The proxy computer can send them its own

copies. For various reasons, proxies are becoming less popular nowadays, but larger organizations might still use them.

You'll need to speak to your system administrator to see if your office uses a proxy. If it does, your administrator will most likely give you an address, which may take the form of a web address or an IP address. Once you have this information, follow these steps to configure the proxy:

1. Open a GNOME Terminal window (Applications ► Accessories ► Terminal).
2. At the prompt, type the following to open the `bash.bashrc` file in Gedit:

```
sudo gedit /etc/bash.bashrc
```

3. Scroll to the bottom and add the following to the file:

```
export http_proxy=http://username:password@address:port_number/
```

Replace *username* and *password* with your own details. If your proxy doesn't use usernames and passwords, simply leave them out (type `http_proxy=http://address:port_number/`). Replace *address* and *port_number* with the relevant details.

4. Add the line shown in step 3 to the `/etc/username/.bashrc` file in exactly the same way. Open the file in Gedit and add the line to the bottom of the file.
5. Save the file, and then log out and back in again.

Tip Some ISPs run proxy servers, too. However, unlike proxies in offices, it's normally up to you whether you choose to use them. You might find that using a proxy speeds up your connection, especially when you access popular sites, so it's worth trying out. To find out if your ISP offers a proxy, visit its technical support web pages or phone its technical support line.

Setting Up Online Software Repositories

Although you installed Ubuntu from the CD-ROM supplied with this book, you should download any additional software from the online Ubuntu software repositories. It might take a few minutes for the software to download, but you're almost certainly guaranteed the latest version of the program in question.

In addition, once the online software repositories are set up, your system will automatically upgrade any software you have to the newest supported versions. This makes for a more secure system, because many software updates fix critical bugs that can leave your system open to illicit interests.

The installation, removal, and updating of software under Ubuntu is handled by the Synaptic Package Manager. This is a little like the Add/Remove Programs applet within the Windows Control Panel, except that it will automatically search for, download, and install any new software you require. As is the Linux way, there's no charge for new software. What more could you ask?

Note *Package* is Linux lingo for a complete file that contains everything you need to install a piece of software on your system.

You can find the Synaptic Package Manager under the System ► Administration menu. Depending on whether the Ubuntu installation program was able to get online, you might find that the Synaptic Package Manager is already set up to use the online repositories. Whatever the case, you should follow these steps to make sure:

1. Select System ► Administration ► Synaptic Package Manager. Adding and removing software makes system-wide changes, so you'll need to enter your password.
2. In the Synaptic Package Manager window, click Settings ► Repositories.
3. Click the Settings button. In the dialog box that appears, make sure there's a check in the Show Disabled Software Sources box, and then click OK.
4. In the Software Preferences dialog box, shown in Figure 8-8, make sure there are checks in the boxes alongside the following entries:

CD Ubuntu 5.10 "Breezy Badger" (Binary)
Officially supported
Restricted copyright

Ubuntu 5.10 "Breezy Badger" (Binary)
Officially supported
Restricted copyright

Ubuntu 5.10 Updates (Binary)
Officially supported
Restricted copyright

Ubuntu 5.10 "Breezy Badger" (Binary)
Community maintained (Universe)

Ubuntu 5.10 Security Updates (Binary)
Officially supported
Restricted copyright

Ubuntu 5.10 Security Updates (Binary)
Community maintained (Universe)

Caution When clicking the Software Preferences dialog box check boxes, be careful *not* to click the check box for any entry that contains the word *Source*. This will activate the source code repositories and will therefore cause confusion when you download software.




Figure 8-8. Configuring the Synaptic Package Manager allows you to grab the latest software from online repositories.

5. Select the following entry:

Ubuntu 5.10 "Breezy Badger" (Binary)
Community maintained (Universe)

6. Click Edit.

7. In the Sections field of the dialog box that appears, add the word `multiverse`, so that the line now reads `universe multiverse`. Then click OK.

8. In the Software Preferences dialog box, click OK.

9. You'll be told that you must now download the latest package lists from the servers. Agree to do so.

After you set up the Synaptic Package Manager to use the online repositories, you will be able to search for software by clicking the Search button. To install a new software package, click the check box alongside it in the list of results, and then click Mark for Installation.

Almost as soon as you download the package lists from the server, you will find a new icon in your notification area (the equivalent of the system tray in Windows) and a speech bubble telling you that system updates are available.

Updating your system software is covered in Chapter 9, and you'll find a full introduction to the way Ubuntu handles software installation and removal in Chapter 29.

ADDITIONAL NOTEBOOK CONFIGURATION

Generally speaking, a notebook computer will not need any configuration above and beyond what's outlined in this chapter. For example, if you have a wireless PCMCIA card, you can simply follow the instructions under the "Configuring a Wireless Card" heading. However, you may also have a touchpad or be able to adjust your CPU's frequency.

Touchpad Configuration

If your notebook has a Synaptics TouchPad, you may want to configure it to your liking. Many notebooks come equipped with Synaptics hardware, although there are also similar touchpads made by Alps and other manufacturers.

If your touchpad is a genuine Synaptics pad, Ubuntu will have installed the correct driver automatically. This should let you use the right edge of the pad to fast scroll through documents and web pages. Sadly, Ubuntu doesn't come with any way of configuring the touchpad (the Mouse applet under the System ► Preferences menu will have no effect). This means that you will be unable to adjust the sensitivity of the pad to your tastes.

However, you can download and use GSynaptics, which allows you to set touchpad preferences. Go to <http://sourceforge.jp/projects/gsynaptics>. Look for the download file created for Ubuntu systems and download it (the file will have a .deb file extension). Then open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type `sudo dpkg -i filename`, replacing *filename* with the name of the file you downloaded.

Before you can use GSynaptics, you need to add a line to your X.org configuration file. You can do this by editing the file in Gedit. Type `sudo gedit /etc/X11/xorg.conf` in the GNOME Terminal window to open the configuration file. Look through the document for the line that reads `Identifier "Synaptics Touchpad"` and then, beneath it, insert a new line that reads `Option "SHMConfig" "true"`. Then save the file and reboot your system. Once Ubuntu is up and running, click System ► Preferences ► Touchpad to start GSynaptics.

If you have an Alps touchpad, you can configure its tap-to-click function using the `tpconfig` utility, which you can download via the Synaptic Package Manager. To find out how to use it, open a GNOME Terminal window and type `man tpconfig`. If all you want to do is turn off tap-to-click behavior, type `sudo tpconfig -to`.

CPU Frequency Scaling

Another handy notebook add-on is the GNOME CPU Frequency Scaling Monitor. If you have a compatible CPU in your notebook, this tool lets you adjust the speed of the chip to save battery power. Most modern mobile-oriented CPUs support this function. Unfortunately, because of its ability to slow down your system, the applet is considered a security risk. Before you use it, you must reconfigure your system to allow it to work.

Open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type `sudo dpkg-reconfigure gnome-applets`. You'll then be asked if you want to set the SUID of the frequency applet. Select Yes. Reboot your computer, and then right-click a blank spot on the Panel at the top of the screen. Click Add to Panel and then, in the dialog box that appears, scroll down to the System & Hardware heading. Click the CPU Frequency Scaling Monitor icon and click the Add button. To alter your CPU frequency, click the applet and choose the clock speed setting you desire.

Configuring E-Mail and Instant Messaging

Being online is all about staying in touch, and Ubuntu is no slouch in this regard. Ubuntu offers a full-featured e-mail program, called Evolution, as well as an instant messaging client called Gaim.

Unlike similar instant messaging clients, Gaim supports the variety of Internet chat protocols, such as ICQ, MSN, Yahoo, and IRC. This means you can chat with friends and colleagues on different networks using this one program.

Evolution is able to work with both IMAP and the popular POP3 mail servers offered by ISPs and used within corporate environments. Additionally, it can work with the Microsoft Exchange protocol used by offices running the Outlook mail program and also Novell GroupWise. We'll look at using Evolution in Chapter 27. Here, you'll learn how to configure the e-mail client to receive and send mail.

Configuring E-Mail Access

Before starting, you'll need to find out the addresses of the mail servers you intend to use. In the case of POP3 and IMAP mail accounts, you'll need to know the incoming and outgoing server addresses (outgoing may be referred to as SMTP). In the case of Microsoft Exchange, you'll need to know the OWA URL and, optionally, the Active Directory/Global Address List server. With Novell GroupWare, you'll simply need to know the server name.

You'll also need to know your username and password details for the incoming and possibly outgoing mail servers.

After gathering the necessary information, follow these steps to configure Evolution:

1. Start the Evolution e-mail client by clicking its icon at the top of the screen, to the right of the menus. Alternatively, you can select Applications ► Office ► Evolution.
2. When Evolution starts for the first time, you'll be invited to enter your configuration details via a wizard. The first screen will ask for your name and the e-mail address you wish to use within Evolution. These are what will appear in outgoing messages. Beneath this is a check box that you should leave checked if you want the account you're about to create to be the default account. In nearly all situations, this will be the correct choice. You can also fill in the Reply-To and Organization information if you wish, but these fields can be left blank. They're not normally displayed by most e-mail clients. Click the Forward button to continue.
3. The next screen asks for details of the receiving (incoming) mail server that you want to use, as shown in Figure 8-9. First, select the server type from the drop-down list. If you don't know which option to go with, select POP. This is by far the most common type of incoming mail server currently in use.
4. In the Configuration section, enter the server address and username in the relevant fields. It's very likely that you'll need to select Always from the Use Secure Connection drop-down list, too. Most POP servers employ at least a password system, and some employ more elaborate protection. You can find out what system your mail server uses by clicking the Check for Supported Types button. Click Forward to continue.



Figure 8-9. Evolution can work with a variety of mail servers, including POP3, Microsoft Exchange, and IMAP.

5. You might need to enter your mail password, depending on which server type you chose. In some cases, you'll need to type this later when you download your mail for the first time. Click Forward to continue.
6. You're given the chance to choose between various additional options, such as how often you want Evolution to check for new mail or if you want to delete mail from the server after it has been downloaded. Unless you have been told otherwise or have special requirements, it should be okay to leave the default settings as they are. If you use a Microsoft Exchange server, you may need to enter the Active Directory/Global Address List server details here. Click Forward to continue.
7. Depending on the server type you chose, you might now need to fill in the outgoing (SMTP) server address. Type this into the Server field. If your SMTP server requires authentication, put a check in the relevant box, and then enter your username. Again, unless you've been specifically told to use this option, you can leave it at its default (inactive) state. Click Forward to continue.
8. You're invited to enter a name for the account. This is the account name you will see when you use Evolution. The default is your e-mail address, but you can type something more memorable if you wish. Click Forward to continue.

9. Finally, choose your location, which will have the effect of automatically defining your time zone. This will ensure that e-mail messages are correctly time-stamped. Click Forward to continue, and then click the Apply button to finish the wizard.

Setting Up Instant Messaging

Instant messaging is a way of chatting with other people in real time. It's as if you were having a phone conversation, but you're typing instead of speaking. You can talk to one other person or a whole group of people, and sometimes share files with them.

The instant messaging program under Ubuntu, Gaim, offers the same functions and works in an almost identical way to programs that you might have used under Windows. It supports virtually all the popular chat standards, such as ICQ/AOL and MSN (Hotmail/Passport). It assumes that you already have an account with each service, which will likely be the case if you've used instant messaging programs under Windows. You can have as many accounts as you wish and can select the one you want to use when you log in.

To transfer your instant messaging account over to Gaim, you just need your screen name and password. As with other instant messaging clients, you'll be able to choose an on-screen alias.

Tip For any queries about how Gaim works, see the FAQ at <http://gaim.sourceforge.net/faq.php>. If you're interested in learning everything there is to know about Gaim, consider reading *Open Source Messaging Application Development: Building and Extending Gaim*, written by Gaim project leader Sean Egan (1-59059-467-3; Apress, 2005).

Follow these steps to set up Gaim:

1. Start Gaim by clicking Applications ► Internet ► Gaim Internet Messenger. When the program starts for the first time, it will automatically open the Add Account dialog box, although it might be behind the main login window. If so, click to bring it to the front of the desktop.
2. In the Add Account dialog box, select the account type you want to set up from the Protocol drop-down list.
3. Enter your screen name, password, and alias details, as required.
4. If you don't want to type your password each time you run Gaim, check Remember Password. However, be aware that someone else using the computer could abuse your account.
5. You can put a check in Auto-Login if you want to log in to your account as soon as Gaim is launched. However, as with many other instant messaging clients, Gaim adds an icon to the notification area, and you can click this to go online and offline.
6. You can put a check in the New Mail Notifications box if you want to be notified of any mail sent to you via the address registered with your instant messaging service.

7. If you want to use a buddy icon (the icon that others will see when they connect to you), click the Open button and browse to a picture.
8. If you wish to connect to a specific instant messaging server or if your network uses a proxy, click Show More Options and enter the details accordingly. In most cases, you won't need to do this.
9. When you've finished, click the Save button. Then, in the main Gaim login window, click Sign On.

After this, you should find the program works just like any other instant messaging program. You can double-click each contact in your list to start a conversation. To sign off, right-click the icon in the notification area and select the option from the menu, as shown in Figure 8-10. To add another account, click the Accounts button in the login window, and then click the Add button in the Accounts window. Then follow the preceding step-by-step instructions.

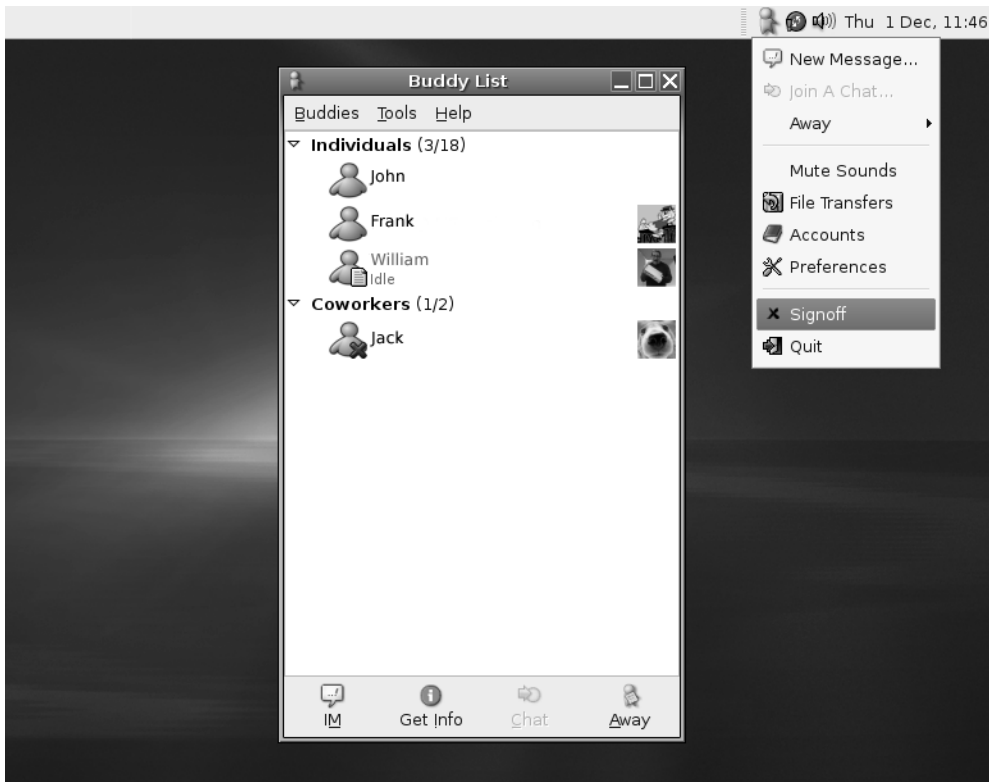


Figure 8-10. You can control Gaim by right-clicking its notification area icon.

Adding a Printer

Most people have a printer nowadays, and Ubuntu supports a wide variety of models—everything from laser printers to color ink-jet models, and even some of the very old dot-matrix printers.

If you work in an office environment, you'll probably be expected to access a shared printer. Sharing a printer is usually achieved by connecting the device directly to the network. The printer itself normally has special built-in hardware to allow this to happen. Alternatively, the printer might be plugged into a Windows computer, such as a Windows NT, 2000, or XP server (or even simply someone's desktop PC), and shared so that other users can access it. Ubuntu will work with network printers of both types.

Configuring a Local Printer

A *local printer* is one that's directly connected to your computer, normally via USB, although if the printer is a number of years old, it might connect via the parallel port. To set up a local printer, follow these instructions:

1. Click System ► Administration ► Printing. In the Printers window, double-click the New Printer icon.
2. In the Add a Printer dialog box, ensure Local Printer is selected. With any luck, your printer will have been automatically detected. But if this isn't the case, click the Printer Port drop-down list and select USB Printer #1 for a USB printer, or Parallel Port #1 for a parallel-port based printer. Click the Forward button to continue.
3. Select your printer manufacturer from the drop-down list, and the printer model from the list below. The printer will be referred to by its full title, rather than just its model number. This includes any prefixes, such as *Optra* or *Stylus Color*. If your printer isn't listed, look for the next best match. For example, my Epson Stylus Color 3600 wasn't in the list, but the Epson Stylus Color 3200 was, so I was able to select this driver. (I knew it would be compatible because the 3200 is similar to the 3600, and they also share the same Windows driver file.) Click Finish.

Tip If you can't find your printer, or any mention of it, try visiting the manufacturer's web site and looking on its support pages for a Postscript Printer Driver (PPD) file. You might also take a look at www.linuxprinting.org/download/PPD/ and www.adobe.com/products/printerdrivers/winppd.html, which offer many printer drivers available for download. If you find one, download it, and then click the Install Driver button to install it.

4. Once installation has finished, the printer will then appear in the Printers window, as shown in Figure 8-11. To see whether it's working correctly, right-click its icon, select Properties, and then click the Print a Test Page button.



Figure 8-11. Once the printer is configured, it appears in the Printers window.

If the printer is installed correctly, you should find yourself with a test page showing color gradations.

If the printer hasn't been installed correctly, it either won't work at all or will start spewing out page after page of junk text. If this is the case, turn off the printer, and then double-click the printer icon. In the window that appears, right-click the printer job and select Cancel. You have either selected the wrong port or, more likely, installed incompatible printer drivers. Right-click the new printer's icon and select Remove. Then repeat the installation steps, this time trying different settings.

Configuring a Network Printer

A network printer is one that is not directly connected to any computer. Instead, it connects to the network via an Ethernet cable. In this way, all computers in the office will be able to use it. If the printer is directly connected to a computer, it will very probably be shared via Windows/SMB. In this case, follow the instructions in the next section.

Some printers have the required server hardware built in, but others might use a special print server module that attaches to the printer's USB or parallel printer port. Ubuntu can work with both types of hardware.

Ubuntu is compatible with Unix (LPD), HP JetDirect, and Internet Printing Protocol (IPP) server types. These are the most ubiquitous types currently in use.

Before beginning, you'll need to find out the printer's network address and, if relevant, the queue name or the port number. You should be able to find out these details by speaking to your network administrator or the person who configured the printer.

Follow these steps to configure a network printer:

1. Click System ► Administration ► Printing. In the Printers window, double-click the New Printer icon.
2. In the Add a Printer dialog box, click Network Printer.
3. In the drop-down list, select the type of printer server. If you're unsure of which to choose, try CUPS Printer (IPP). If you wish to connect to a Hewlett Packard printer, select HP JetDirect.

4. Enter the network address of the printer. In the case of HP JetDirect, the default port number should be okay, unless you have been specifically told to enter a different number. Depending on which server option you chose, you may also need to enter the queue name.
5. Choose the printer make and model (see step 3 in the previous section for guidance). Then click Finish.
6. When the printer is installed, right-click its icon in the Printers window, select Properties, and then click Print a Test Page.

If the printer doesn't work, it's likely that you set the wrong server type. Try an alternative type; if you chose CUPS the first time, try HP JetDirect the second time. Many print servers can emulate a variety of modes, so trying a different setting may work.

If the printer starts spewing out page after page of text, it's likely that you selected an incorrect printer driver. Cancel the job at the printer. Then double-click the printer icon, right-click the job, and select Cancel. Then right-click the printer, select Remove, and repeat the installation steps, this time trying an alternative driver.

Configuring a Windows/SMB Shared Printer

A Windows (or SMB) printer is one that's directly connected to a computer, and then made available across the network. Effectively, the computer acts as the printer server. Often, in corporate environments, such printers are attached to server computers, but an individual may share the printer attached to a workstation.

In a home situation, a Windows/SMB share is an excellent and inexpensive way of sharing a printer among many computers. The printer is attached to one PC, and, as long as that computer is switched on, the printer will be available to the other computers in the household.

Assuming that the printer has been correctly set up to be shared on the host computer, connecting to a Windows/SMB printer share is easy. In most cases, Ubuntu will do the hard work for you.

Follow these steps to set up a Windows/SMB shared printer:

1. Click System ► Administration ► Printing. In the Printers window, double-click the New Printer icon.
2. In the Add a Printer dialog box, click Network Printer.
3. In the drop-down list, select Windows Printer (SMB), as shown in Figure 8-12.
4. Wait for a minute or two while Ubuntu probes the network to see if any shared printers are available. When Ubuntu discovers a printer, you might find an Authentication dialog box appears. Here, you should type the username and password (if applicable) for the shared printer. Ask the person who administers the computer that's sharing the printer for this information. (When connecting to the shared printer attached to a Windows XP Home machine, I was able to click the Cancel button in the Authentication dialog box.)

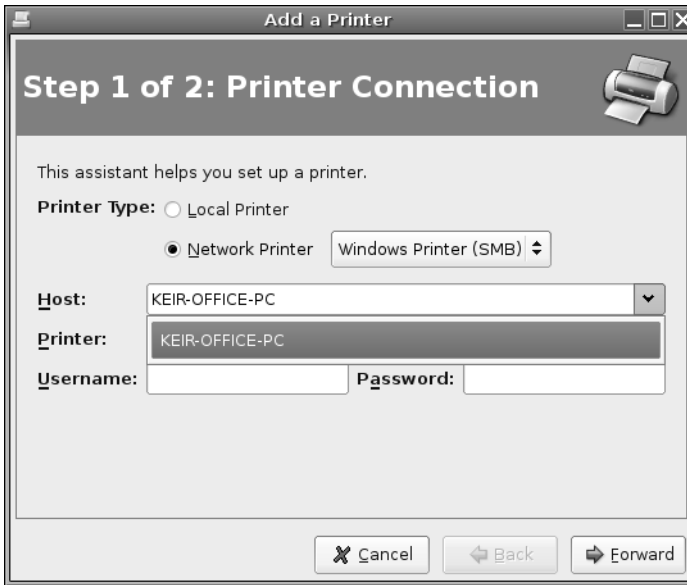


Figure 8-12. *Ubuntu should be able to automatically detect any Windows or SMB shared printers on your network.*

5. In the Host drop-down list, select the computer to which the printer is attached. If there is more than one shared printer on the network, you may need to choose between several options, so make sure that you know the correct network name of the machine sharing the printer. The best way to do this is to speak to your system administrator or the person who set up the printer share.
6. Wait a few more seconds, and then, in the Printer drop-down list, select the shared printer. Once again this should be detected automatically.

Note If you find that the Host and Printer details aren't automatically configured, you will need to enter them manually. Speak to the system administrator or the individual in charge of the shared printer to find out what these are.

7. In the Username and Password fields, type the username and password that might be required to access the shared printer, if they haven't already been filled in automatically. (With the Windows XP Home host I used for testing, I found that I could type any username and leave the Password field blank.) To connect to a Windows XP Professional or Windows 2000 host, you will most likely need to enter both a username and password. These can be the login details of any user of the computer or, if the shared computer and printer are configured for Guest access, you can try typing Guest for the username and leaving the Password field blank.

8. Select the printer driver (see step 3 in the “Configuring a Local Printer” section for guidance). Then click Finish.
9. When the printer is installed, right-click its icon in the Printers window, select Properties, and then click Print a Test Page.

If the printer doesn't work at all, the username and password details may be wrong. Alternatively, if you entered the printer network address manually, you might have entered it incorrectly. Try repeating the installation steps and using alternative settings.

If the printer makes a noise as if to start printing, but then decides not to, you might need to change a setting *on the Windows machine*. Click Start ► Settings ► Printers and Faxes, and then right-click the shared printer's icon. Select Properties, and then click the Ports tab in the window that appears. Remove the check on the Enable Bidirectional Support box and click OK. Then restart both the Windows and Ubuntu computers.

If the printer starts spewing out page after page of text instead of the test page, it's likely that you selected an incorrect printer driver. Cancel the job at the printer. Then double-click the printer icon in Ubuntu, right-click the job, and select Cancel. Next, right-click the printer and select Remove. Then repeat the installation steps, this time trying an alternative driver.

Setting Up Digital Imaging Devices

One of the ironies of computers is that they're ruthlessly logical yet allow people to pursue their creative hobbies. Photography is just one example. If you want to transfer your photos into your Ubuntu system, you can use one of the two usual ways: get them from your digital camera or scan them in using a scanner. We'll look at both methods here, starting with digital cameras.

Transferring Photos from a Digital Camera

Digital cameras have been around for a while now and offer a genuine alternative to traditional film photography. They store pictures on computer memory cards rather than on film, meaning their images can quickly and easily be downloaded to a PC.

You can transfer your pictures from your digital camera to your computer via a direct cable connection into your USB or serial port, or by using a card reader. The latter requires buying an extra piece of hardware into which you insert the memory card from the camera so you can download images from it. Ubuntu supports both ways of transferring digital images.

Note Very nearly all card readers are supported by Ubuntu. If you find that your camera doesn't work via a direct cable connection, consider buying a card reader.

Using a Card Reader

When you use a card reader, Ubuntu simply mounts the contents of the memory card into a folder on your hard drive. I'll describe how mounting works in Chapter 14, but basically it makes a virtual folder on your hard disk, from which you can access the contents of the memory card. In most cases, you can read, delete, and even write new contents to the card in this way.

No extra software is necessary, and you can simply use Nautilus, the standard hard disk file browser under Ubuntu, to access the folder's contents.

Note If there's no card in the reader, it probably won't be identified by Ubuntu when you plug it in. The reader will not be recognized until a card is present.

With a card reader, the setup procedure is simple: attach it to your PC, and then insert the memory card. Within seconds of doing so, you should see an icon for your card appear on the desktop. You'll also find that *all* the card reader slots are accessible via the Places ► Computer menu, even if they don't have any cards in them, as shown in Figure 8-13. This can prove handy if, for whatever reason, the desktop icon doesn't appear.

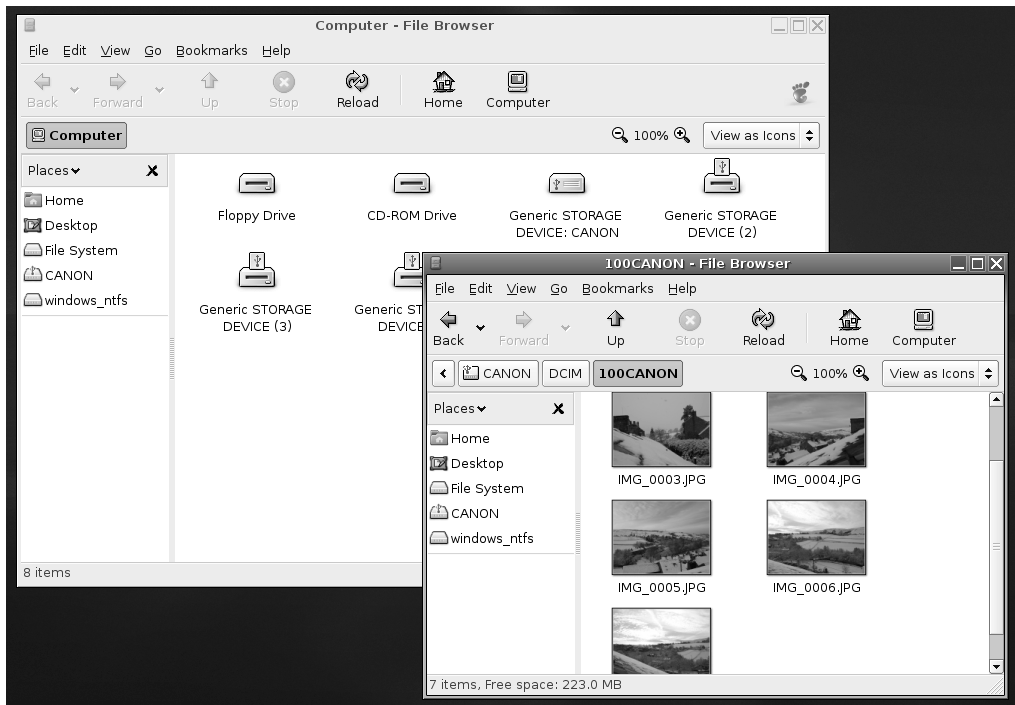


Figure 8-13. Whenever you insert a memory card, it will be made available in the Computer file browser window, as well as an icon on the desktop.

Double-click the desktop icon to open the virtual folder containing the card contents, and you can then copy them to folders on your hard disk or otherwise manipulate the files.

A very important rule must be followed when accessing memory cards with Ubuntu: the card must be *unmounted* before you remove it from the card reader. This is quite simple to do.

Just right-click the desktop icon and select Unmount Volume. Make sure that you save and close any files that you may have been working on before you unmount the card.

Following this, you can safely remove the card. Reinserting it will make it available once again.

Caution Be very careful not to remove a memory card from a card reader while you're writing or reading from it on your PC. This will most likely damage the card irreparably. At the very least, it will wipe the contents of the card, so you'll lose your photographs.

Using a Camera's Cable Connection

Setting up a digital camera cable connection is usually easy. In most cases, you can simply plug it in to your USB port, turn it on, and Ubuntu will recognize it and make its contents available in a folder (in a virtually identical way to that described in the previous section for a card reader).

Ubuntu comes complete with the gThumb software, which includes support for the vast majority of older cameras (those a few years old that directly connect to your PC by its serial or USB port). Because of the way these cameras work, they are not usually automatically configured by Ubuntu, hence the need for a separate software package.

To configure your camera through gThumb, follow these steps:

1. Select Applications ► Graphics ► gThumb Image Viewer.
2. Click File ► Import. If you've already set up your camera in the program, skip to step 5.
3. The first time you use the program, you'll need to set up your camera. To do so, click the camera icon at the top left. In the Camera Model dialog box, select your model from the drop-down list. This should fill in the Port field automatically, but it's a good idea to check that this setting is correct. If your camera isn't listed, select the nearest match. (For example, my Olympus C920 Zoom wasn't listed, but the Olympus C900 Zoom choice proved to be compatible.) Click OK.
4. Close the Import Photos dialog box. Then open it again by clicking File ► Import. This time, your camera will be recognized, and thumbnails of the photos stored on the camera will be shown in the dialog box, as shown in Figure 8-14.
5. In the Import Photos dialog box, select the photos that you wish to import by clicking them (hold down the Ctrl key to select multiple photos). Then click the Import button. The pictures will transfer across to your PC. Bear in mind that this might take some time if you're using a serial connection! It might appear as if the program is doing nothing, even though it is in fact working.
6. After importing has finished, you should find your pictures available in your /home directory. When gThumb imports the pictures, it creates a directory based on the date and time.

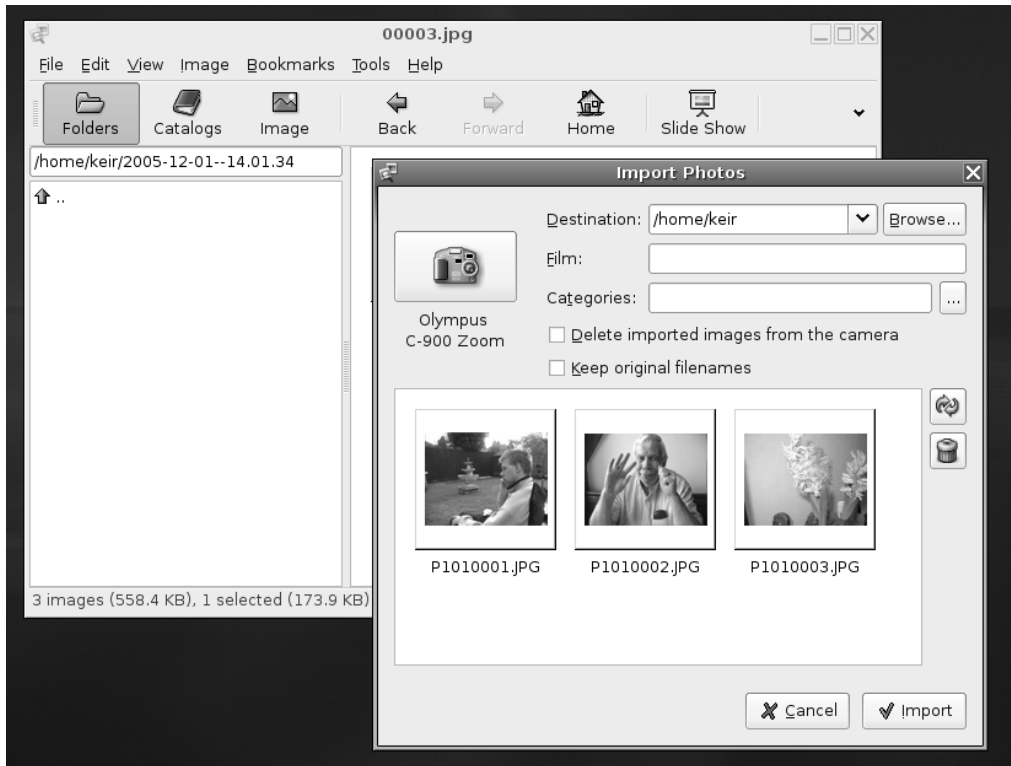


Figure 8-14. *gThumb Image Viewer supports most of the older-generation cameras that connect directly to your computer.*

Configuring a Scanner

Although scanners have fallen out of favor recently with the advent of digital photography, they're vital for getting nondigital photos and old documents onto your PC.

A lot of flatbed scanners can be made to work under Ubuntu, but not all types are supported. You can check the list of currently supported scanners by visiting www.sane-project.org/sane-mfgs.html. Additional models are added to the list all the time, and this is another reason to make sure your system is completely up-to-date (see Chapter 9).

The best test of whether your scanner is supported under Ubuntu is simply to test if it will work. Scanning within Ubuntu is handled by the XSane utility. This is a stand-alone program that operates like the TWAIN drivers that you might have used under Windows, except it's a lot more powerful.

Tip XSane is even capable of optical character recognition! Simply use the Synaptic Package Manager to download and install the `gocr` program. Then select the OCR button on the toolbar of XSane's image preview window.

To configure a scanner, and also scan images, follow these steps:

1. Select Applications ► Graphics ► XSane Image Scanning Program. On startup, the program will attempt to detect your scanner. If it finds a compatible model, you will then see its findings in a splash screen, in which you can click the OK button. If the scanner hasn't been detected, you'll be told that, and XSane won't start.
2. The main XSane program window is similar in appearance to the TWAIN scanner drivers you might have used under Windows, as shown in Figure 8-15. At the top of the window is the XSane mode drop-down list. Here, you can select from a variety of scanning modes, such as those to scan documents for faxing. However, in most cases, the Viewer setting is best. This lets you preview your scans and then save them to disk if you're happy with them.

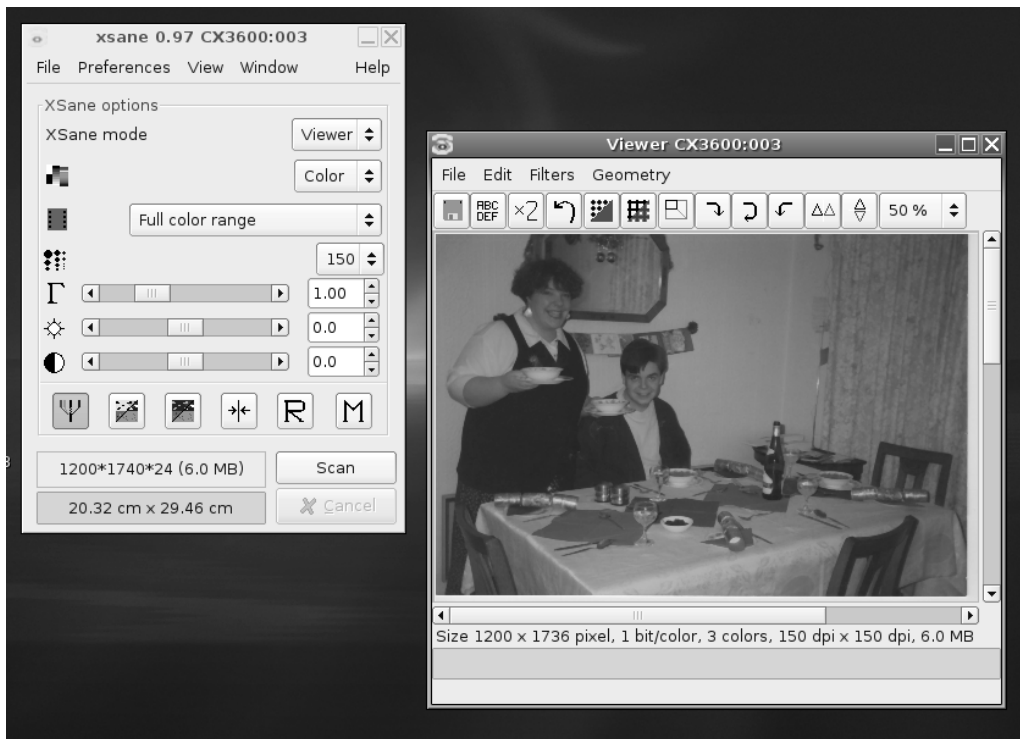


Figure 8-15. The XSane program works a little like TWAIN drivers under Windows and makes scanning easy.

3. Beneath this is the color/monochrome drop-down list. Here, you can select to scan Binary (line-art), Gray (grayscale), or Color.
4. Below this, you'll find the drop-down list that sets the scanning type. You should select Full Color Range for scanning photographs or artwork, but you might also select slide (transparency) and negative scanning modes, if you have a compatible scanner.

5. Next is the dots per inch (DPI) setting. Generally speaking, 300 DPI is acceptable for scanned photos, while 150 DPI will be acceptable for artwork such as diagrams.
6. Beneath are the gamma, brightness, and contrast sliders, in that order. You can use these sliders to adjust the quality of the scan.
7. To scan a preview, click the Scan button.
8. When the scan is complete, the image viewer window will display the scan (note that XSane is clever enough to crop the picture and delete any space around it). Here, you can choose to manipulate the image by rotating it or applying cleaning filters (look under the Filters menu). Once you're happy with the results, click the Save button on the preview window's toolbar (the first button on the left).
9. If you wish, you can adjust the brightness/contrast settings and scan again. Just close the original preview and click the Scan button.

Using a USB Memory Stick

Although at one time the floppy disk drive ruled when it came to transferring small files between computers, nowadays, the USB memory stick has taken its place. These small devices, often incorporated into key rings, contain nonvolatile memory and retain their contents even when no power is applied.

The good news is that Ubuntu works with just about every make and model of memory stick. Simply insert the device, and then wait a few seconds while Ubuntu recognizes it. It will then make the contents of the stick available via a desktop icon. Alternatively, you can click Places ► Computer to access the device. The standard Nautilus file browser window is used to access the device, so you can copy, delete, and rename files as usual.

An important caveat is that when you're finished with the device, you *shouldn't* simply unplug it. Instead, you need to *unmount* it first. Make sure you've saved and closed any files on the memory stick. You might also need to close any Nautilus file browser windows that are browsing the stick. Then right-click its desktop icon and select Unmount Volume.

Installing 3D Graphics Card Drivers

Virtually all graphics cards are automatically supported and configured within Ubuntu, so you can stick with the default X.org drivers installed by Ubuntu. However, if you want to use their 3D functionality—usually to play 3D games or run 3D modeling software—then some extra steps may be necessary.

ATI and Nvidia 3D cards are well supported under Ubuntu, including the Radeon series and the GeForce series.

Adding 3D graphics support involves downloading and installing special driver software from the Ubuntu repositories via the Synaptic Package Manager. If you haven't yet configured the Synaptic Package Manager, see the "Setting Up Online Software Repositories" section earlier in this chapter. The drivers are provided by the manufacturer of the graphics card and are closed source (sometimes referred to as *binary only* or *proprietary*).

Note If you experience seemingly random system-wide crashes or freezing after installing a 3D graphics driver, consider reverting to your old setup by using the Synaptic Package Manager to simply remove the old driver. If you find your system won't boot to the GUI after removing the drivers, type `sudo dpkg-reconfigure xserver-xorg` at the command prompt to reconfigure your graphics settings.

Installing an ATI Driver

To install support for ATI 3D cards, follow these steps:

1. Select System ► Administration ► Synaptic Package Manager.
2. Click the Search button and enter `xorg-driver-fglrx` as a search term. Find your card among those listed in the Description box. Click the box next to the entry in the results list and select Mark for Installation. Then click Apply.
3. When the Synaptic Package Manager has finished installing the driver, open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type `sudo fglrxconfig`. This will configure your X.org configuration file to work with the ATI drivers. Note that you will need to answer some questions about your keyboard and mouse. In most cases, the default answers should work fine.
4. Once the configuration program has finished, reboot your system.
5. When the system is up and running, you will be able to further configure the card by opening a GNOME Terminal window and typing `fireflycontrolpanel`.

There are mixed reports about the effectiveness of the ATI drivers, and some people report they don't work very well. If you find that the drivers crash X, or prevent you from running a GUI, try manually configuring the X server to use the new drivers by typing `sudo dpkg-reconfigure xserver-xorg`. When prompted to choose a driver for your graphics card, select `fglrx` from the list. Complete the X.org setup wizard, and then restart your computer.

If you find the drivers *still* don't work, then it's time to give up and restore the original settings: simply run `sudo dpkg-reconfigure xserver-xorg` and choose either the ATI driver or, if that doesn't work, the Vesa driver.

Installing a Nvidia Driver

The Nvidia driver supports practically all Nvidia 3D cards produced in recent times—everything from the latest GeForce cards to the oldest TNT cards from the mid-1990s.

To install support for Nvidia cards, follow these steps:

1. Select System ► Administration ► Synaptic Package Manager.
2. Click the Search button and enter `nvidia-glx` as a search term. In the list of results, click the check box next to `nvidia-glx` and also `nvidia-settings`, so that both are marked for installation. Then click Apply.

3. After the installation has completed, open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type `sudo nvidia-glx-config enable`.
4. Reboot your system.
5. To further configure the Nvidia card once your PC is back up and running, open a GNOME Terminal window and type `nvidia-settings`. You'll see a Settings window like the one shown in Figure 8-16.

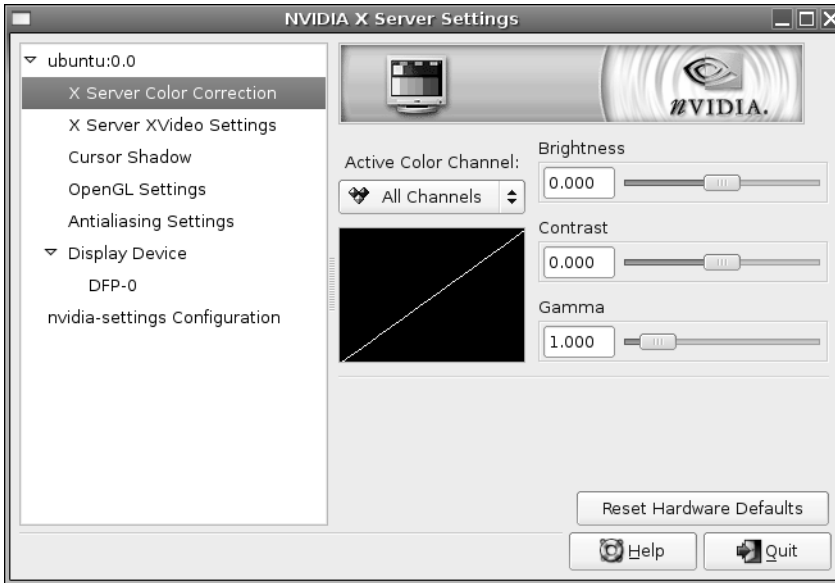


Figure 8-16. Once the correct driver has been installed, you can configure Nvidia cards by typing `nvidia-settings` in a GNOME Terminal window.

Testing 3D Capabilities

To test the 3D function of the graphics card, try running a screensaver. If the graphics drivers are correctly installed, the animations should run smoothly and with a high frame rate.

To select the screensaver, click System ► Preferences, and then click the Screensaver icon. Select an option from the list that is preceded by *GL*, although you might also try some of the Ant screensavers, such as AntSpotlight.

Click the Preview button to see the screensaver in action. Note that these screensavers are visually very impressive. This is a good chance to show off your new Ubuntu setup to friends and colleagues!

Configuring Bluetooth

Bluetooth is the short-range networking facility that allows various items of hardware, usually those designed for mobile devices, to work with each other wirelessly. You can use Bluetooth for everything from file transfers between a mobile phone and computer to employing a wireless keyboard with your desktop computer.

For Bluetooth to work, both devices need to have Bluetooth support. Many mobile phones come with Bluetooth nowadays, and an increasing number of notebook computers do, too. It's also possible to buy very inexpensive Bluetooth dongles that attach to the USB port of your computer.

Bluetooth hardware is automatically recognized under Ubuntu, and the low-level driver software is installed by default. Therefore, all you normally need to do is install the software that provides the Bluetooth functionality you require.

Pairing Bluetooth Devices

When two pieces of Bluetooth-compatible hardware need to communicate on a regular basis, they can *pair* together. This means that they trust each other, so you don't need to authorize every attempt at communication between the devices.

Pairing is very simple in practice and works on the principle of a shared personal ID number (PIN). The first Bluetooth device generates the PIN, and then asks the second Bluetooth device to confirm it. Once the user has typed in the PIN, the devices are paired.

Pairing is easily accomplished under Ubuntu and doesn't require any additional software. As an example, I paired my Ubuntu test PC to a Nokia 6680 mobile phone. It's easiest to initiate pairing on the phone, which should then autosense the PC's Bluetooth connection. On the Nokia 6680, I opened the menu and selected Connections ► Bluetooth. Then I hit the right arrow key to select Paired Devices and selected Options ► New Paired Device ► More Devices. This made the phone autosense my Ubuntu PC, which was identified as ubuntu-0. I was then prompted to enter a PIN.

The default PIN for pairing under Ubuntu is 1234, but you can change this by opening a GNOME Terminal window (Applications ► Accessories ► Terminal) and typing the following, which will open the PIN configuration file in Gedit:

```
sudo gedit /etc/bluetooth/pin
```

In the configuration file, replace 1234 with the PIN you want to use. Then save the file. Next, type the following in the GNOME Terminal window:

```
sudo /etc/init.d/bluez-utils restart
```

This will restart Ubuntu's Bluetooth service to take into account your reconfiguration.

Some users report that they are unable to pair their phone with Ubuntu unless they edit another configuration file. If you find that pairing is refused, type this command:

```
sudo gedit /etc/bluetooth/hcid.conf
```

This will open the Bluetooth configuration file in the Gedit text editor. Look for the line that reads as follows:

```
pin_helper /usr/bin/bluez-pin;
```

And change it so it reads like this:

```
pin_helper /usr/bin/bluepin;
```

Save the file and restart the Bluetooth service:

```
sudo /etc/init.d/bluez-utils restart
```

Then try again to pair the device and your PC.

Transferring Files Between Bluetooth Devices

If you own a Bluetooth-equipped camera phone, you might be used to transferring pictures to your computer using Bluetooth. It's by far the easiest way of getting pictures off the phone and avoids the need for USB cables or card readers.

In order to transfer files via Bluetooth, you'll need to install some additional software from the Ubuntu repositories. If you haven't yet configured the Synaptic Package Manager, see the "Setting Up Online Software Repositories" section earlier in this chapter. Then open Synaptic Package Manager (from the System ► Administration menu). Click the Search button and enter `gnome-bluetooth` as a search term. In the list of results, click the check box alongside the entry, and then click Mark for Installation.

You should now find two new entries on the Applications ► System Tools menu: Bluetooth File Sharing and Bluetooth Manager. I found that Bluetooth Manager didn't work very well with my phone, so I will ignore it here. It's designed to allow you to manage Bluetooth connections, and you may like to explore its functions later.

The following instructions detail how to transfer any kind of file to and from your PC using Bluetooth. Once again, I use a Nokia 6680 in the examples, but the instructions should work with any phone, or even any Bluetooth device capable of sending and receiving files.

Note Some phones refuse to transfer files unless the phone and computer are paired, so follow the instructions in the previous section first. Phones like the Nokia 6680 don't need pairing for file transfer, although each transfer will need to be confirmed each time.

Sending Files to a Ubuntu PC

Follow these steps to send files from a Bluetooth device to your PC:

1. Select Applications ► System Tools ► Bluetooth File Sharing. Nothing will appear to have happened but, in fact, a new icon will have been added to the notification area. This indicates your computer is ready for incoming Bluetooth connections.
2. On the Bluetooth device from which you wish to send the file, start the file transfer. On the Nokia 6680, I clicked the file and selected Send ► Via Bluetooth.
3. When the file transfer is initiated, a dialog box will appear on your computer asking if you wish to accept the file, as shown in Figure 8-17. Click OK. (If the two devices are paired, the file transfer may happen instantly without the confirmation dialog box.) The file will be saved to your `/home` directory.



Figure 8-17. If you send files from a Bluetooth device to your PC, you may be asked to authorize receipt.

Sending Files from a Ubuntu PC to Another Device

The easiest way to send files from your PC to a Bluetooth device is to create a desktop shortcut onto which you can drag-and-drop files. Follow these steps to create the shortcut:

1. Right-click the desktop and click Create Launcher.
2. In the Name field, type something like **Send file via Bluetooth**.
3. In the Command field, type `gnome-obex-send`.
4. You can also choose to give the new shortcut an appropriate icon. Click the icon button, and then type the following into the Path field:

```
/usr/share/icons/hicolor/48x48/stock/io/stock_bluetooth.png
```

5. Click OK.

After you've created the icon, you can send files as follows:

1. Drag-and-drop a file onto the launcher (icon) you just created.
2. The Choose Bluetooth Device dialog box appears, as shown in Figure 8-18. Click Refresh to make the computer detect any nearby Bluetooth devices. Remember that your device will need to be set to be "visible" so that other Bluetooth devices can automatically detect it.
3. Select the device to which you want to transfer the file, and then click OK.
4. Check the device to see if the file transfer needs to be authorized. If the devices are paired, the transfer might take place automatically.

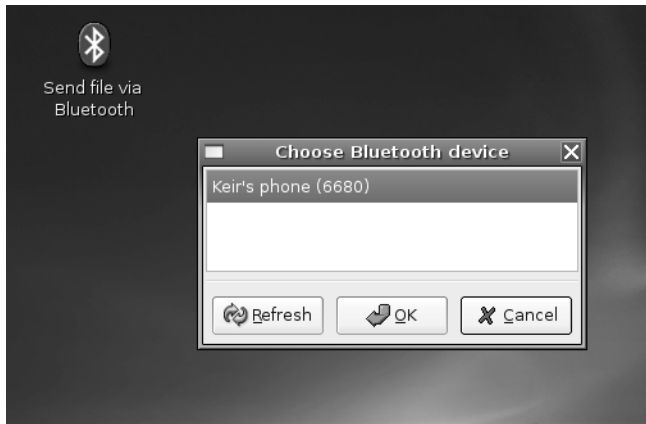


Figure 8-18. Sending files from your Ubuntu PC is easy if you create a desktop shortcut.

Using a Bluetooth Keyboard or Mouse

You may find that your Bluetooth-equipped keyboard or mouse works automatically under Ubuntu. However, if not, you may find the following instructions useful:

1. Open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type `hcitool scan`.
2. Your Bluetooth keyboard or mouse should be identified in the results (ignore any other devices that might appear in the list). If not, make sure it isn't in sleep mode. You might also have to press a button on the device for it to be made visible.
3. Alongside the entry for the keyboard or mouse will be a MAC address—a series of numbers like `00:12:62:A5:60:F7`.
4. In the GNOME Terminal window, type the following:

```
sudo hidd --connect xx:xx:xx:xx:xx:xx
```

Replace `xx:xx:xx:xx:xx:xx` with the series of numbers you discovered in the previous step.

5. You should now find that your keyboard or mouse works under Ubuntu. You now need to make sure your mouse or keyboard works every time you boot your computer, so you'll need to edit the Ubuntu Bluetooth configuration file. Type the following in the GNOME Terminal window to open the configuration file in Gedit:

```
sudo gedit /etc/default/bluez-utils
```

6. Search for the line that reads `HIDD_ENABLED=0` and change it to `HIDD_ENABLED=1`.
7. Beneath this will be a line that begins `HIDD_OPTIONS=`. Change this so it reads like this:

```
HIDD_OPTIONS="--connect xx:xx:xx:xx:xx:xx --server"
```

Once again, `xx:xx:xx:xx:xx:xx` is the MAC number you discovered earlier.

8. If you find that the keyboard or mouse doesn't connect upon reboot, try step 7 again, but this time, change the `HIDD_OPTIONS` line so it reads like this:

```
HIDD_OPTIONS="-i xx:xx:xx:xx:xx:xx --server"
```

Again replacing `xx:xx:xx:xx:xx:xx` with the MAC address you discovered earlier.

9. Save the file and reboot to see if the mouse or keyboard is working.

Tip If you want to quickly connect a Bluetooth keyboard or mouse to your computer, but don't need to make it permanent, just open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type `sudo hidd --search`.

Configuring Sound Cards

Generally speaking, your sound card shouldn't require any additional configuration and should work immediately after you install Ubuntu. The icon for the volume control applet is located at the top right of the Ubuntu desktop, and it offers a quick way to control the master volume.

However, if your sound card offers more than stereo output, such as multiple-speaker surround sound, then it's necessary to take some simple steps to allow full control of the hardware:

1. Right-click the volume control icon (the one that looks like a speaker) and select Open Volume Control.
2. In the dialog box that appears, click Edit, and then click Preferences.
3. The Volume Control Preferences dialog box appears, as shown in Figure 8-19. Select the sliders that you wish to be visible. For example, on my desktop computer that has 5.1 surround sound, I was able to add a slider for the center and back speakers. On my notebook that has a sound card featuring pseudo-surround sound, I was able to add a control to alter the intensity of the effect.
4. When you've finished, click the Close button.

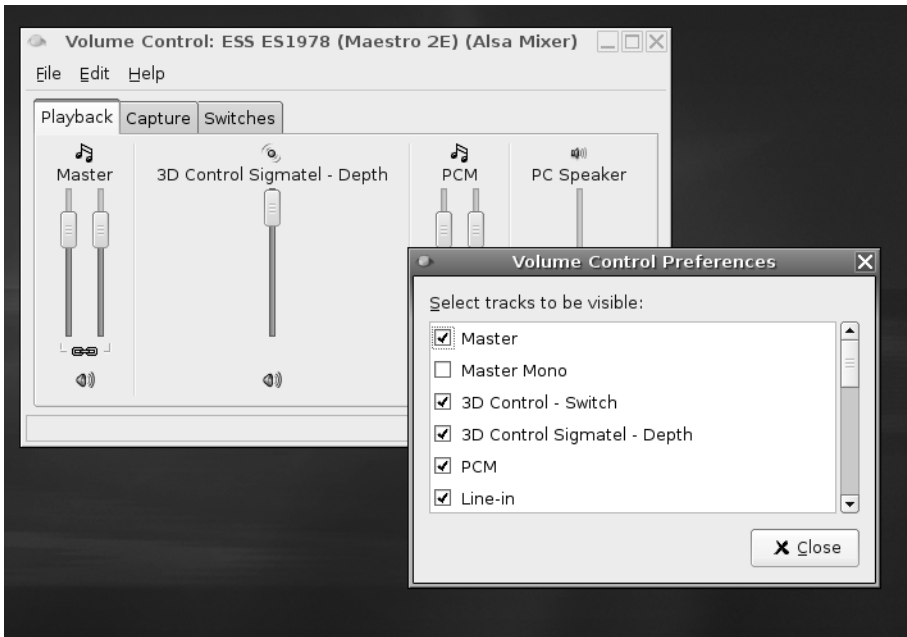


Figure 8-19. You can add sliders to control all aspects of your sound card's output.

Summary

In this chapter, you learned how to set up just about every piece of hardware you might have attached to your computer. Additionally, we looked at configuring various software components within your Ubuntu setup that are vital for its correct functioning.

We stepped through getting online with Ubuntu (including joining a wireless network), configuring e-mail, adding a printer, setting up online software repositories, setting up a digital camera, configuring a 3D graphics card, and much more.

In Chapter 9, we move on to look at how you can ensure that your system is secure and protected from hackers.



How to Secure Your Computer

Linux is widely considered to be one of the most secure operating systems around. On a basic level, Linux is built from the ground up to be fundamentally sound, and it forces users to work with security in mind. For instance, it enforces the system of ordinary users who are limited in what they can do, thus making it harder for virus infections to occur.

In addition, Linux contains a firewall that is hard-wired into the kernel. It's called iptables (www.netfilter.org) and is considered among the best by practically all computer security experts. Not only that, but it can protect your home PC just as well as it can protect the most powerful supercomputer. But, as with many Linux kernel components, iptables is difficult to use. It requires in-depth knowledge of how networks operate and an ability to hack configuration files, both of which are beyond the skills of many ordinary computer users. Fortunately, several programs act as interfaces to iptables and make it simple to operate (or at least as simple as any equivalent Windows-based software firewall, such as Zone Labs' ZoneAlarm).

Perhaps surprisingly, Ubuntu doesn't install any firewall configuration program by default. According to the official FAQ (www.ubuntulinux.org/support/documentation/faq), the developers consider Ubuntu to have no need for such a thing. However, configuring the firewall with a program like Firestarter, which we examine later in this chapter, can be done so quickly and with such little effort that there's no reason *not* to make use of the Linux firewall.

In this chapter, you'll learn how to configure the Linux firewall, but first, you'll spend some time examining more basic security concepts. Following that, we'll look at some elementary steps that you can take to protect your system.

Windows Security vs. Linux Security

If you've switched to Ubuntu from Windows, there's a very good chance that the security failings of Windows featured in your decision. By any measure, Microsoft's record on security within its products is appalling. A new and serious security warning appears seemingly on an ongoing basis, and a new and devastating virus makes news headlines with similar frequency (usually described as "a PC virus" rather than what it actually is: a Windows virus).

One argument is that Windows is the target of so many viruses merely because it's so popular. Although it's true that some of the underground crackers who write viruses dislike Microsoft, there's also little doubt that Windows has more than its fair share of security issues.

The situation is certainly getting better but, even so, Microsoft's latest operating system, Windows XP, provides many good examples of why it's an easy target. Upon installation, the default user is given root powers. True, a handful of tasks can be performed only by the genuine administrator, but the default user can configure hardware, remove system software, and even wipe every file from the hard disk, if he pleases. Of course, you would never intentionally damage your own system, but computer attackers use various techniques to get you to run malicious software (by pretending it's a different file, for example) or by simply infecting your computer across the Internet without your knowledge, which is how most worms work.

Viruses and worms also usually take advantage of security holes within Windows software. As just one example, a famous security hole within Outlook Express allowed a program attached to an e-mail message to run when the user simply clicked a particular message to view it. In other words, infecting a Windows machine was as easy as sending someone an e-mail message!

It's a different story with Linux. Viruses and worms are far rarer than they are on Windows. In fact, the total number of viruses and worms that have been found in the wild infecting Linux systems number far less than 100 (one report published in 2003 put the number at 40, and the number is unlikely to have grown much since then). Compare that to Windows, where according to the Sophos antivirus labs (www.sophos.com), approximately 1,000 new viruses are discovered *every month*! The Sophos antivirus product now guards against just under 100,000 viruses.

Note The high number of Windows viruses may be due to the quantity of Windows PCs out there. After all, for a virus to spread, it needs computers to infect, and it won't have trouble finding other Windows computers.

But while I would love to say that security holes are not found on Linux, the sad truth is that they're a fact of life for users of every operating system. Many so-called rootkits are available, generated by members of underground cracking groups. These are specialized software toolkits that aim to exploit holes within the Linux operating system and its software.

The bottom line is that while writing a virus or worm for Linux is much harder than doing the same thing on Windows, all Linux users should spend time defending their system and *never* assume that they're safe.

Root and Ordinary Users

As I've mentioned in earlier chapters, Linux makes use of something called the *root* user account. This is sometimes referred to as the *superuser* account, and that gives you an idea of its purpose in life: the root user has unrestricted access to all aspects of the system. The root user can delete, modify, or view any file, as well as alter hardware settings.

Linux systems also have ordinary user accounts, which are limited in what they can do. Such users are limited to saving files in their own directory within the /home directory (although the system is usually configured so that an ordinary user can read files outside the /home directory, too). But an ordinary Ubuntu user cannot delete or modify files other than those that he created or for which he has explicitly been given permission to modify by someone else.

On most Linux systems, it's possible to type `root` at the login prompt and, after providing the correct password, actually log in as root and perform system maintenance tasks. Ubuntu is slightly different in that the root account is disabled by default, and users are instead able to borrow superuser powers whenever they're required. For this to happen, they need to provide their login password. With desktop programs, this is automatic, but at the command prompt, users need to preface commands with `sudo`.

Although the root account is disabled, most key operating system files “belong” to root, which is to say that only someone with superuser powers can alter them. Ordinary users are simply unable to modify or delete these system files, as shown in Figure 9-1. This is a powerful method of protecting the operating system configuration from accidental or even deliberate damage.

Note Along with the root and ordinary user accounts, there is a third type of Linux account, which is similar to a limited user account, except that it's used by the system for various tasks. These user accounts are usually invisible to ordinary users and work in the background. For example, the audio subsystem has its own user account that Ubuntu uses to access the audio hardware. The concepts of users and files are discussed in more depth in Chapter 14.

ARE YOU A CRACKER OR A HACKER?

Linux users are often described as *hackers*. This doesn't mean they maliciously break into computers or write viruses. It's simply using the word *hacker* in its original sense from the 1970s, when it described a computer enthusiast who was interested in exploring the capabilities of computers. Many of the people behind multinational computing corporations started out as hackers. Examples are Steve Wozniak, a cofounder of Apple Computer, and Bill Joy, cofounder of Sun Microsystems.

The word *hacker* is believed to derive from model train enthusiasts who “hacked” train tracks together as part of their hobby. When computing became popular in the early 1970s, several of these enthusiasts also became interested in computing, and the term was carried across with them.

However, in recent years, the media has subverted the term *hacker* to apply to an individual who breaks into computer systems. This was based on ignorance, and many true hackers find the comparison extremely offensive. Because of this, the term *cracker* was invented to clearly define an individual who maliciously attacks computers.

So, don't worry if an acquaintance describes herself as a Linux hacker, or tells you that she has spent the night “hacking.” Many Linux types use the term as a badge of honor.

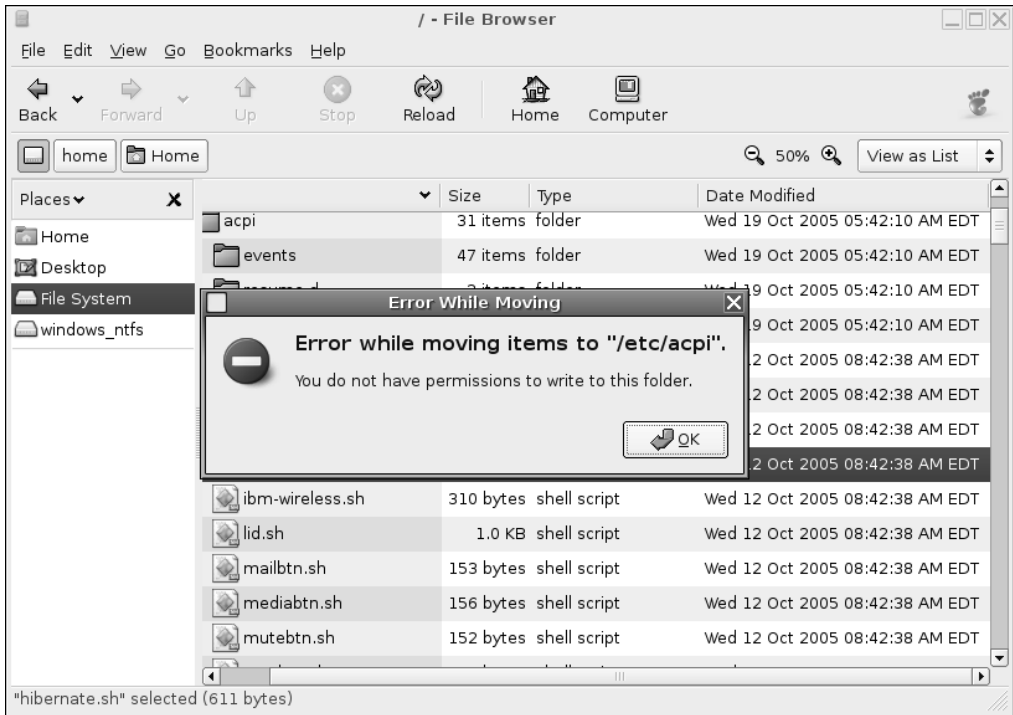


Figure 9-1. Ordinary users are simply unable to modify or delete essential system files under Linux.

Common-Sense Security

As you start to understand how Ubuntu works, you'll become more and more aware of common-sense methods that will protect your system. However, I'll outline a few of these now to get you started:

Entering your password: Be very wary if you're asked to enter your password (outside of initial login, of course). You'll be asked to provide your password when following many of the configuration steps within this book, for example, and this is acceptable and safe. But if you're asked to do so out of the blue, then you should be suspicious. If the root password prompt dialog box (shown in Figure 9-2) appears when you run a file that shouldn't really need root permissions, such as an MP3 or OpenOffice.org file, you should treat the situation with caution.



Figure 9-2. Beware if you're asked to type your password out of the blue and for no apparent reason.

Installing new software: Be careful in choosing programs to download and install. Because Linux works on the basis of open-source code, anyone can theoretically tamper with a program, and then offer it for download by the unwary. This very rarely happens in real life. Even so, it's wise to avoid downloading programs from unofficial sources, such as web sites you find online via a search engine and whose authenticity you cannot totally trust. Instead, get software from the web site of the people who made it in the first place or, ideally, from the official Ubuntu software repositories (discussed in Chapter 8).

Updating your system: Always ensure your system software is completely up-to-date. As with Windows, many Ubuntu programs have bugs that lead to security holes. Crackers target such vulnerabilities. Downloading the latest versions of Ubuntu software ensures that you not only get the latest features, but also that any critical security holes are patched. As with most versions of Linux, updating Ubuntu is easy and, of course, it's also free of charge. You'll learn how to get online updates in the next section.

Locking up your PC: Limit who has physical access to your computer. Any Ubuntu system can be compromised by a simple floppy boot disk, or even the Ubuntu installation CD. Booting a PC using such disks gives anyone complete root access to your system's files, with no limitations. This is for obvious reasons; the idea of a boot disk is to let you fix your PC should something go wrong, and you cannot do this if you're blocked from accessing certain files. When Linux is used on servers that hold confidential data, it's not uncommon for the floppy and CD-ROM drives to be removed, thus avoiding booting via a boot disk. Such computers are also usually locked away in a room or even in a cupboard, denying physical access to the machine.

WHERE'S THE ANTIVIRUS?

At first glance, it may appear that there are very few Linux antivirus programs. Actually, many of these exist, but they're designed to work on server computers and primarily guard against Windows viruses, in addition to the handful of Linux viruses. The idea is that they protect Windows users who access the server.

Very few antivirus products are aimed at the Linux desktop. However, one example includes F-Secure's Anti-Virus for Linux Workstations. This costs around 80 euros (just under \$100) and is available from www.f-secure.com/estore/fsavlinuxwks.shtml. AVG (www.grisoft.com) and Kaspersky (www.kaspersky.com) also produce Linux workstation versions of their antivirus products.

The main issue with all of these programs is that they're not open source, as with most of the Linux software included in Ubuntu. If you absolutely must have your entire system running free software, consider ClamAV (www.clamav.net). This is a product designed to work on Linux servers but is flexible enough to run on desktop computers, too. ClamAV is included in the Ubuntu software repository, and so is available via the Synaptic Package Manager program. Be aware that ClamAV is a command-line program, however. You'll need to read its man page to learn how it works. In addition, you might choose to read the online documentation at www.clamav.net/doc.

Online Updates

The Ubuntu notification area (the equivalent of the Windows system tray) contains a program that automatically monitors the package repositories and tells you when updates are available. This is an extension of the Synaptic Package Manager program called Update Manager. If you've already configured the Synaptic Package Manager, as described in Chapter 8, and haven't yet updated your system, this icon will have probably turned into a red circle with two arrows in it. This is informing you that updates are available. In addition, each time you boot, you will see a speech bubble telling you that updates are available. When your system is completely up-to-date, the icon will disappear.

Clicking the Update Manager icon opens the Software Updates window, as shown in Figure 9-3. To go online and grab the updated files, click the Install button at the bottom-right side of the window. You will probably be asked to enter your root password, because system files will need to be altered.




Figure 9-3. You'll be informed if your system is in need of updates, and the Update Manager can take care of everything for you.

Be aware that some updates can be large and might take some time to download, particularly if you're doing it for the first time after installing Ubuntu. (On my system, I had to download almost 200MB of files on my first update.)

Once the downloads have finished, you probably won't need to reboot unless the kernel file has been updated. The Software Update program will inform you if you need to restart the computer.

The Ubuntu Firewall

A *firewall* is a set of programs that protects your PC when it's online. It does this by watching what data attempts to enter your PC from the Internet and allowing in only what it is sure is secure (which usually is what you've asked for). It also attempts to close off various aspects of your Internet connection, so that crackers don't have a way in should they target your system.

Although Ubuntu includes a powerful firewall in the form of iptables, you'll also need a program that can manage it. Here, I'll show you how to use Firestarter, available from the Ubuntu software repository, for this purpose. Together with the built-in firewall, this really does provide industrial-level protection.

The benefit of configuring the firewall is that even if your system has security vulnerabilities because of buggy software, crackers will find it a lot harder to exploit them across the Internet. When someone attempts to probe your system, it will appear to be virtually invisible.

Caution Although software firewalls such as the one built into Linux offer a high level of protection, it's best to use them in concert with a hardware firewall, such as that provided by most DSL/cable broadband routers (curiously, some of these routers actually use Linux's iptables software as well). Many security experts agree that relying solely on a software firewall to protect a PC affords less than the best level of protection.

Installing Firestarter

Let's get started by downloading and installing Firestarter. Follow these steps:

1. Select System ► Administration ► Synaptic Package Manager. Click the Search button and enter `firestarter` as a search term. In the list of results, locate the program and click the check box. Then choose to install the package.
2. After installing Firestarter, log out and then back in again (to update the menus to show Firestarter).
3. Once the desktop is back up and running, select Applications ► System Tools ► Firestarter. When you run Firestarter for the first time, it will walk you through a wizard.
4. Click to continue the wizard beyond the introductory page.
5. The first step asks which network interface Firestarter should configure, as shown in Figure 9-4. If you use an Ethernet card, have a wireless card, or attach a broadband modem directly to your computer, the answer will probably be `eth0` or `wlan0`. However, if you use a modem, the answer is `ppp0`.



Figure 9-4. Firestarter includes a wizard to walk you through the basics of firewall configuration.

6. Put a check in the IP address is assigned via DHCP box, *unless* you're using a modem.
7. You're asked if you want to enable Internet connection sharing. This allows you to turn your computer into an Internet router and can be very useful in certain circumstances. You can activate this later on by running the wizard again (to rerun the wizard, simply click Firewall on Firestarter's main window, and then click Run Wizard).
8. Save your settings. The Firestarter main window then opens.

Configuring Firestarter

Firestarter works by controlling the data that goes in and out of your computer via your Internet or network connection. By default, it blocks every type of uninvited inbound connection but allows every type of outbound connection. This needs some explanation.

Whenever you click a link on a web page, your computer sends a request for data to the web server hosting the web page. Within a few milliseconds, that data will be sent to your computer. This is an *inbound* data connection. The Linux firewall is clever enough to realize that the data was requested by you, so it is allowed through. However, any uninvited connections are turned away. If, out of the blue, someone attempts to connect to your computer via the popular Secure Shell (SSH) tool, as just one example, he won't be allowed to make that connection. This is a good thing because it makes your computer secure. Crackers are turned away whenever they try to connect, no matter *how* they try to connect.

But in some circumstances, allowing uninvited connections is useful. For example, if you create a shared folder for other computers in your office to connect to, they will frequently make uninvited inbound connections to your computer. And if you want to make use of SSH to connect to your computer remotely, you will need to allow such incoming connections. Therefore, Firestarter lets you allow through certain types of inbound connections.

Outbound traffic is any kind of data originating on your computer that is sent out on the network and/or Internet. By default, Firestarter allows out all data, no matter what it is. This is described as a *permissive* policy. But Firestarter can be configured to block all outgoing connections *apart from those you opt to allow through*. This is described as a *restrictive* policy and can be useful in blocking certain types of programs that “phone home” with personal data about you, such as spyware. It can also prevent certain types of viruses and worms from spreading. The downside is that you must configure Firestarter to take into account every type of outgoing data connection, such as those for web browsers, instant messaging programs, and so on.

You can configure Firestarter by clicking the Policy tab in the main program window. Click the Editing drop-down list and choose to configure either the inbound traffic policy or the outbound traffic policy.

Note Firestarter is used only to configure the built-in firewall and doesn't need to be running for the firewall to work. Once you've finished configuration, you can quit the program. You'll need to use it again only if you wish to reconfigure the firewall.

Setting Inbound Rules

For most users, Firestarter's default inbound traffic policy will be perfectly acceptable. It configures the firewall to disallow all uninvited incoming data connection, apart from certain diagnostic tools, such as ping, traceroute, and so on. You can choose to disallow those as well, as described shortly in the “Turning Off Diagnostic Services” section.

You might wish to allow an incoming connection if you intend to connect to your computer via SSH from a remote location or if you have a shared folder created for other computers in your office. It's a must if you're running the BitTorrent file sharing application. Additionally, if you run a web, e-mail, or other type of server on your computer, you will need to allow the correct type of incoming connection here.

Here's how to set inbound connection rules:

1. In the Firestarter main window, click the Policy tab. Select Inbound Traffic Policy in the Editing drop-down list.
2. Right-click in the second box on the Policy tab (with the headings Allow Service / Port / For), and then select Add Rule.
3. The Add New Inbound Rule dialog box appears. In the Name drop-down list, select the type of outgoing connection you want to allow, as shown in Figure 9-5. To allow others to access shared folders on your computer, select Samba (SMB). To allow SSH or BitTorrent connections to your computer, select the relevant entry from the list. Selecting the service will automatically fill in the Port box, which you shouldn't alter unless you know exactly what you're doing.

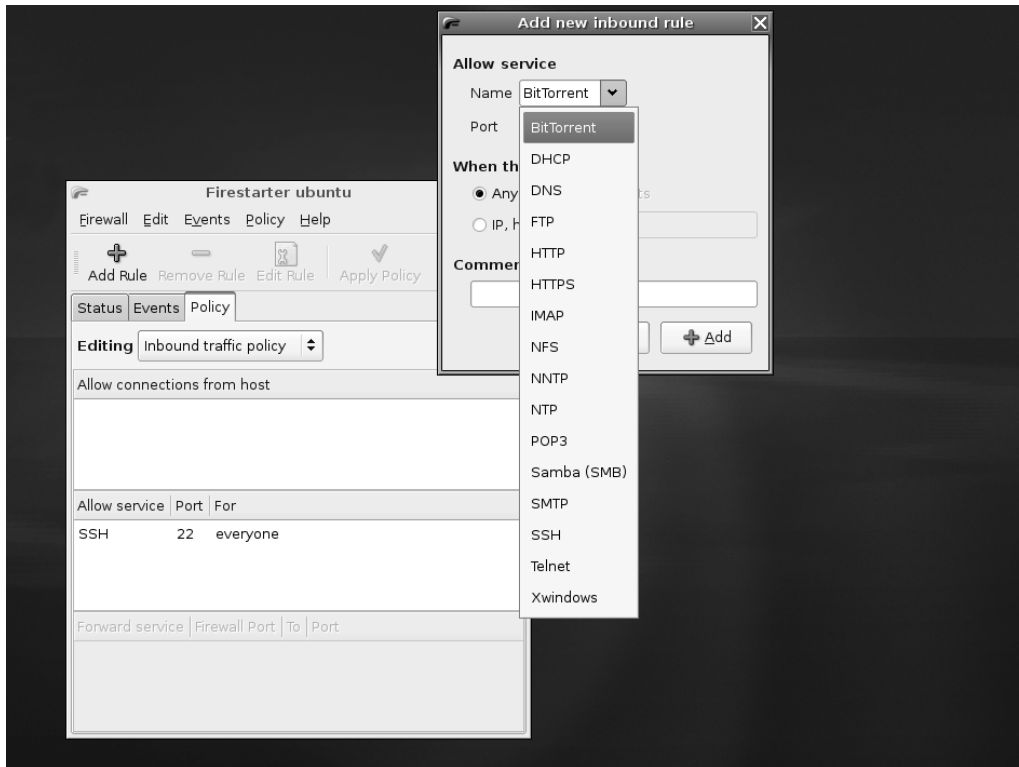


Figure 9-5. *Creating an inbound rule enables computers to connect to your PC uninvited.*

4. If you know the IP address of the computer that's going to make the incoming connection, you can click the **IP, Host or Network** radio button, and then type in that address. However, the default of **Anyone** will allow anyone using any IP address to connect to your computer.
5. Click **Add**. Back in the main Firestarter window, click the **Apply Policy** button.

Note You'll need to return to Firestarter whenever you activate new services on your computer. For example, in Chapter 12, we will look at accessing Windows shares across a network, and you'll need to enable SMB incoming and outgoing access for this to work. In Chapter 34, we will look at using the SSH service, which will have to be allowed through the firewall. In other words, securing your computer isn't something you can do once and then forget about. It's a continual process.

Setting Outbound Rules

By default, Firestarter allows all types of outgoing connections and, as with its incoming connections policy, this is by no means a bad choice for the average user. It's certainly the option that involves the least fuss. However, by opting to go with a restrictive traffic policy, you can completely control what kind of data leaves your computer. Any type of data connection

that isn't authorized will be refused; as far as the program sending the data is concerned, it will be as if your computer did not have a network or Internet connection.

Here's how to set outbound connection rules:

1. In the Firestarter main window, click the Policy tab. Select Outbound Traffic Policy in the Editing drop-down list.
2. Click the Restrictive by Default, Whitelist Traffic radio button.
3. In the second empty box at the bottom of the Policy tab, right-click and select Add Rule.
4. The Add New Outbound Rule dialog box appears. Select the type of data connection you wish to allow. At the very least, you should select HTTP. This will allow your web browser to operate correctly (it's also needed to allow the Synaptic Package Manager and Update Manager programs to work). You should also add a rule for POP3 and another for SMTP, without which your e-mail program won't work. Selecting the type of service will fill in the Port box automatically. You shouldn't alter this unless you know what you're doing.
5. Click the Add button to add the rule. Back in the Firestarter main window, click Apply Policy.
6. Test your settings with a program that uses the services you've just authorized.

Caution If you created an inbound rule, you'll need to create a matching outbound rule. If you created an incoming rule for BitTorrent, for example, you'll need to create an outgoing rule for BitTorrent, too.

You can delete both incoming and outgoing rules by right-clicking their entries in the list.

Turning Off Diagnostic Services

Certain network tools can be misused by crackers in order to break into a computer or just cause it problems. In the past, the traceroute and ping tools, among others, have been used to launch denial-of-service (DoS) attacks against computers.

Ubuntu is set to allow these tools to operate by default. If you want to adopt a belts-and-braces approach to your computer's security, you can opt to disable them. If you don't know what ping and traceroute are, you're clearly not going to miss them, so there will be no harm in disallowing them. Here's how:

1. In the Firestarter main window, click Edit ► Preferences.
2. On the left side of the Preferences window, click ICMP Filtering. Then click the Enable ICMP Filtering check box, as shown in Figure 9-6. *Don't* put a check in any of the boxes underneath, unless you specifically want to *permit* one of the services.

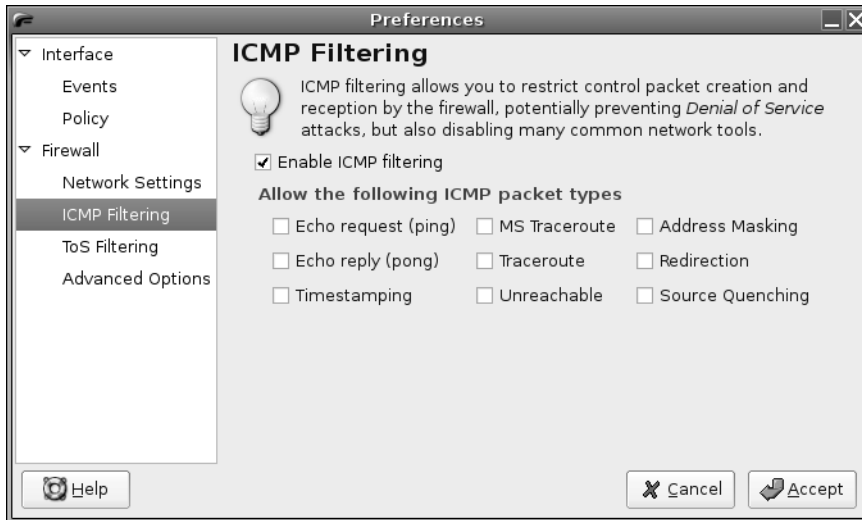


Figure 9-6. By deactivating traceroute, ping, and other services, you can add extra protection to your PC.

3. Click the Accept button to finish.

PARANOIA AND SECURITY

There's a fine line between security and paranoia. Using Firestarter gives you the opportunity to ensure your system is secure, without needing to constantly reassess your system for threats and live in fear.

When considering your system security, remember that most burglars don't enter a house through the front door. Most take advantage of an open window or poor security elsewhere in the house. In other words, when configuring your system's security, you should always select every option and extra layer of security, even if it might not appear to be useful. You should lock every door and close every window, even if you don't think an attacker would ever use them.

Provided a security setting doesn't impact your ordinary use of the computer, you should select it. For example, deactivating the ping response of your computer might sound like a paranoid action, but it's useful on several levels. First, it means your computer is less easy to detect when it's online. Second, and equally important, it means that if there's ever a security flaw in the ping tool (or any software connected with it), you'll be automatically protected.

This illustrates how you must think when configuring your system's security. Try to imagine every situation that might arise. Remember that you can never take too many precautions!

Summary

In this chapter, we've looked at what threats your system faces and how security holes can be exploited by malicious interests. You learned about measures you can take to protect your system, such as updating it online and configuring the system's firewall. We also discussed some common-sense rules you can follow to keep your system safe.

In the next chapter, we move on to looking at how your Ubuntu system can be personalized and how to set up everything to suit your own preferences.



Personalizing Ubuntu: Getting Everything Just Right

If you've read this book from Chapter 1, by this stage, you no doubt have become comfortable with Ubuntu. You've started to realize its advantages and are on the way to making it your operating system of choice.

But things might still not be quite right. For instance, you might find the color scheme is not to your tastes. Or perhaps you feel that the mouse cursor moves a little too fast (or too slowly). Maybe you simply want to stamp your own individuality on your system to make it your very own.

That's what this chapter is all about. We look at personalizing Ubuntu so that you're completely happy with your user experience.

Changing the Look and Feel

Ubuntu is similar to Windows in many ways, but the developers behind it introduced improvements and tweaks that many claim make the software easier to use. For example, Ubuntu offers multiple virtual desktops—long considered a very useful user-interface feature that seems to have passed Microsoft by. It also moves the programs menu to the top of the screen, leaving the whole width of the screen at the bottom to display taskbar buttons. This is very sensible, because the buttons don't look cramped when more than a handful of applications are open. However, if you're not satisfied with Ubuntu's out-of-the-box look and feel, you can change it.

You might be used to changing the desktop colors or wallpaper under Windows, but Ubuntu goes to extremes and lets you alter the look and feel of the entire desktop. Everything from the styling of the program windows to the desktop icons can be altered quickly and easily.

Altering the Theme

Ubuntu refers to the look of the desktop as a *theme*. Because it's built on the GNOME desktop, Ubuntu allows you to radically personalize your desktop theme. Several different themes come with the distribution, and you can download many more themes. Each lets you change the way the windows look, including the buttons and the icon set (although some themes come without additional icons).

However, unlike Windows themes, most GNOME themes don't change the fonts used on the desktop, and the wallpaper and color scheme will probably remain broadly the same. You can change these manually, as described in the “Setting Font Preferences” and “Changing the Wallpaper” sections a bit later in this chapter.

To alter the theme, select System ► Preferences ► Theme. Then it's simply a matter of choosing a theme from the list in the Theme Preferences dialog box, as shown in Figure 10-1. A useful hint is to open a Nautilus file browser window in the background (Places ► Desktop), so you can see how the changes will affect a typical window.

Note The default Ubuntu theme is called Human and is designed to represent the skin tones of the world's population. This is intended to reflect Ubuntu's mission of being accessible to everyone, no matter where or who they are.

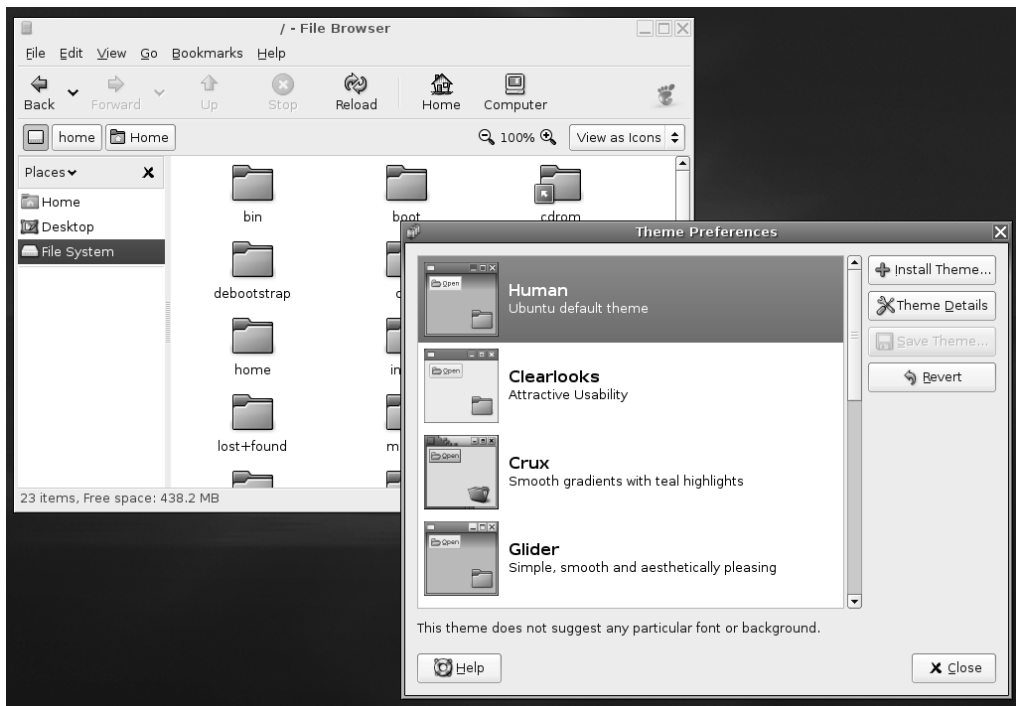


Figure 10-1. Ubuntu comes with several theme choices.

My favorite themes are Clearlooks and Mist, largely because they're simple and uncomplicated. Remember that you'll be working with the theme on a daily basis, so it should be practical and not too distracting. Those miniature close, minimize, and maximize buttons might look stylish, but they're useless if they're so small that you can't reliably click them with your mouse.

As well as changing the overall theme, you can also modify individual theme components, and even download more theme components.

Changing Individual Theme Components

You can alter the three aspects that constitute a GNOME theme: the *controls* (sometimes known as *widgets*), the *window borders*, and the *icons*. Controls are simply the elements you click within dialog boxes and windows: buttons, scroll bars, and so on. The window borders are, as seems obvious, the borders of program windows and dialog boxes, with particular attention paid to the top of the window, where the program name appears along with the minimize, maximize, and close buttons.

Note To make matters a little confusing, some window borders have their own selection of close, minimize, and maximize controls, which can't be overridden with individual selections for controls.

To make changes to a theme, click the Theme Details button in the Theme Preferences dialog box (Figure 10-1), and then click each tab to see your choices, as shown in Figure 10-2. Unfortunately, there are no thumbnail previews of each style, but as soon as you click each option, it will be automatically applied to the currently open windows. To preview the effects fully, the best policy is to keep a Nautilus window open (Places ► Desktop).

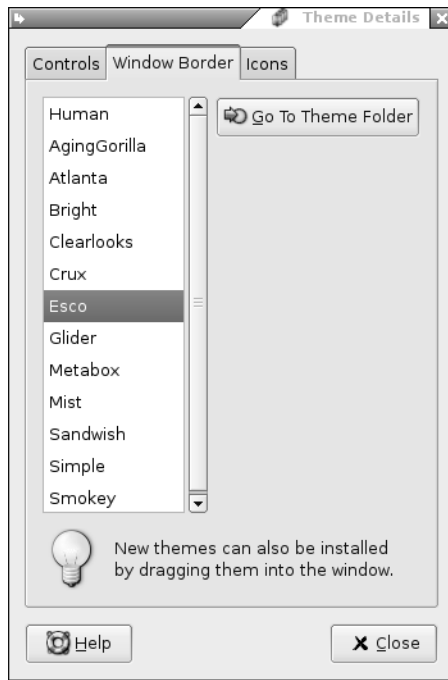


Figure 10-2. You can create a theme by choosing you own controls, window borders, and icons.

When you've made your choices, you can save the theme for further use. Simply click the Save Theme button in the Theme Preferences dialog box. You'll need to give the theme a name

and, if you wish, a short description for future reference. If you don't save the theme, as soon as you select another one, the changes you made will be lost.

Installing Additional Components

If you get tired of the built-in possibilities, you can download additional theme components, such as window borders and controls, to enhance your desktop experience. You have two ways of getting new themes:

- Download themes from the official Ubuntu repositories.
- Visit the GNOME Art web site (<http://art.gnome.org>) and download items from there.

Downloading from Ubuntu Repositories

To get theme components from the Ubuntu software repositories, you use the Synaptic Package Manager. Setting up Synaptic Package Manager to use the online repositories is described in Chapter 8.

Select **System** ► **Administration** ► **Synaptic Package Manager**, click the **Search** button, and enter `gtk2-engines` as a search term (`gtk2-engines` is how Ubuntu refers to theme components). In the list of results will be those `gtk2-engines` already installed, indicated by a dark green check box, and several that are available for download.

Icons rarely come in `gtk2-engines` packages, and instead are contained in their own packages. To find icons, use the Synaptic Package Manager to search for `gnome icon theme` (without any dashes).

Although each theme component comes with a description, you won't really know what it looks like until you see it. The best policy is to download all of them and audition them one by one. However, be aware that themes can be large, so they may take some time to download on a slower connection.

Caution Unless you've already installed the KDE desktop, don't download `gtk2-engines-gtk-qt`. This is a piece of system software designed to give KDE applications the same look as GNOME applications. Selecting it will cause the entire KDE desktop subsystem to download, too.

Don't forget that you're downloading theme components, rather than entire themes. To use your new theme components, select **System** ► **Preferences** ► **Theme**, click the **Theme Details** button, and choose from the various lists.

Downloading from the GNOME Art Web Site

Visiting the GNOME Art site (<http://art.gnome.org>), shown in Figure 10-3, gives you access to just about every theme ever created for GNOME. In fact, the site also contains wallpaper selections, icons, and much more besides. All of the offerings are free to use, and most of the packages are created by enthusiasts.

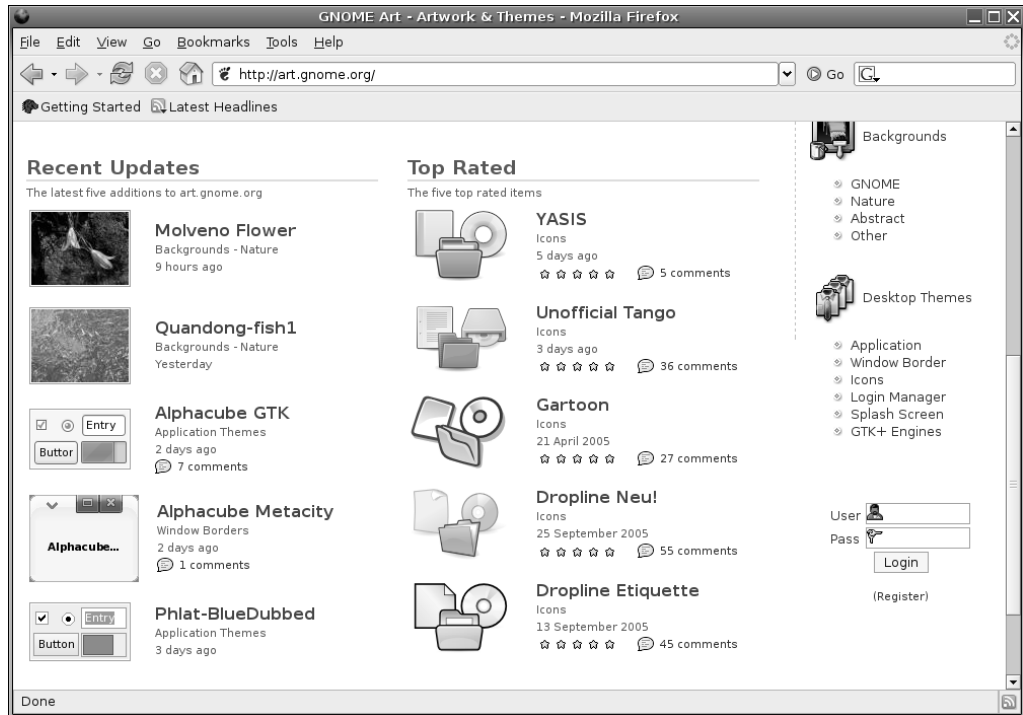


Figure 10-3. The GNOME Art site contains the latest themes, and you can use all of them with Ubuntu.

Installing new theme components is easy. If you wish to install a new window border, for example, click the link to browse the examples and then, when you find one you like, click to download it. It will be contained in a `.tar.gz` archive, but you don't need to unpack it. Simply select **System** ► **Preferences** ► **Theme** and click the **Install Theme** button in the Theme Preferences dialog box. Then browse to the downloaded theme and click **Open**. You can delete the downloaded file when you're finished.

Note The same principle of sharing that underlines the GPL software license is also usually applied to themes. This means that one person can take a theme created by someone else, tweak it, and then release it as a new theme. This ensures constant innovation and improvement.

Changing the Wallpaper

The default Ubuntu wallpaper, Lagoon, is a love-it-or-hate-it affair. Some find its emphasis on dark colors depressing; others appreciate its humanist metaphor. If you're one of those who prefer something different, it's easy to switch. Simply right-click the desktop and click **Change Desktop Background**. If you want to use a picture of your own as wallpaper, click the **Add Wallpaper** button, and then browse to its location.

In the Style drop-down list, you can select from the following choices:

Centered: This option places the wallpaper in the center of the screen. If the wallpaper is not big enough to fill the screen, a border appears around the edge. If it's bigger than the screen, the edges of the wallpaper are cropped off.

Fill Screen: This option forces the picture to fit the screen, including squashing or expanding it if necessary (known as *altering its aspect ratio*). If the wallpaper isn't in the same ratio as the screen, it will look distorted. Most digital camera shots should be okay, because they use the same 4:3 ratio as most monitors (although if you have a wide-screen monitor, a digital camera picture will be stretched horizontally).

Scaled: Like the Fill Screen option, this option enlarges the image if it's too small or shrinks it if it's too big, but it maintains the aspect ratio, thus avoiding distortion. However, if the picture is in a different aspect ratio than the monitor, it may have borders at the edges.

Tiled: If the picture is smaller than the desktop resolution, this option simply repeats the picture (starting from the top left) until the screen is filled. This option is primarily designed for patterned graphics.

Don't forget that the GNOME Art web site (<http://art.gnome.org>) offers many wallpaper packages for download.

Tip Looking for some good wallpaper? Visit Flickr (www.flickr.com). This is a community photography site where many people make their pictures publicly available.

Setting Font Preferences

Ubuntu lets you change the fonts that are used throughout Ubuntu (referred to as *system fonts*). You can also alter how they're displayed.

To change a system font, select System ► Preferences ► Font. In the Font Preferences dialog box, shown in Figure 10-4, click the button next to the system font you want to change, and then choose from the list. You can also set the font point size so, for example, you can make the labels beneath icons easier to read.

By clicking the entries under the Font Rendering heading in the Font Preferences dialog box, you can change how fonts look on your monitor. This will alter the *antialiasing* and *hinting* of the font. Antialiasing softens the edges of each letter to make them appear less jagged. Hinting affects the spacing and shaping of the letters. Used together, they can make the on-screen text look more pleasant. Try each Font Rendering setting in sequence to see which looks best to you (the text in the dialog box will update automatically to show the changes). Nearly everyone with a TFT-based screen, including notebook users, finds the Subpixel Smoothing option best.



Figure 10-4. You can alter the way fonts appear on screen by using the Fonts Preferences dialog box.

BYTECODE HINTING

Two font hinting subsystems are available under Ubuntu: *Autohinting* and *Bytecode Interpreting*. There's a lengthy debate about which produces the best results. Personally, I prefer to use the Bytecode Interpreter because I believe it leads to the cleaner fonts, but others say Autohinting is better in this regard.

Unfortunately, the Bytecode Interpreter is protected by patents in the United States (for more information, see www.freetype.org/patents.html). This means that individuals in the United States can't use it without paying a license fee to the patent holder (in theory, at least). Those in the rest of world should be fine.

It's easy to activate the Bytecode Interpreter. Follow these steps:

1. Open a GNOME Terminal window (Applications ► Accessories ► Terminal).
2. In the terminal window, type the following:


```
sudo dpkg-reconfigure fontconfig
```
3. On the first screen of the configuration program, select Native to activate the Bytecode Interpreter. Alternatively, you can choose Autohinting or even None, which will turn off the hinting system.
4. On the next screen, you can select whether subpixel rendering is activated. This is useful only for TFT screens (including notebooks), so either select Automatic or, if you use a TFT monitor and want to ensure the option is activated, select Always.
5. The third screen offers the option of using bitmap fonts. These are fonts that, unlike the TrueType fonts used within the rest of Ubuntu, don't scale beyond their original size. There's no harm in enabling them because they can sometimes be useful as system fonts.
6. Once the program has finished configuring the software, restart your X server by logging out and then back in again.

Configuring Input Devices

Mouse and key repeat speeds are personal to each user, and you may find the default Ubuntu settings not to your taste, particularly if you have a high-resolution mouse such as a gaming model. Fortunately, changing each setting is easy. You'll find the relevant options under the System ► Preferences menu.

Configuring Mouse Options

Select System ► Preferences ► Mouse to open the Mouse Preferences dialog box, which has three tabs:

Buttons: This tab lets you set whether the mouse is to be used by a left-handed or right-handed person. Effectively, it swaps the functions of the right and left buttons. Beneath this is the double-click timeout setting. This is ideal for people who are less physically dexterous because the double-click speed can be slowed down. On the other hand, if you find yourself accidentally double-clicking items, you can speed it up.

Cursors: On this tab, you can select from any mouse cursor themes that are installed. You can also activate the Locate Pointer option, which causes a box to appear around the mouse cursor when you press the Ctrl key. This can help you find the cursor on a busy desktop.

Motion: This tab, shown in Figure 10-5, lets you alter the speed of the mouse pointer, as well as the drag-and-drop threshold. Changes are made as each setting is adjusted, so to test the new settings, simply move your mouse. Here's what the settings do:

- The Acceleration setting controls how fast the mouse moves. Whenever you move the mouse, the pointer on screen moves a corresponding amount. However, the cursor actually increases in speed the more you move your hand (otherwise, you would need to drag your hand across the desk to get from one side of the screen to the other). This is referred to as *acceleration*. If you set the acceleration too high, the pointer will fly around the screen, seemingly unable to stop. If you set it too slow, you'll need to ramp the mouse several times to make it go anywhere.
- The Sensitivity setting controls how quickly the acceleration kicks in when you first move the mouse. Choosing a higher setting means that you can move the mouse relatively quickly before it starts to accelerate and cover more screen space. A low setting means that acceleration will begin almost as soon as you move the mouse. Higher sensitivity settings give you more control over the mouse, which can be useful if you use image-editing programs, for example.
- The Threshold setting determines the amount of mouse movement allowed in a click-and-drag maneuver before the item under the cursor is moved. This setting is designed for people who have limited dexterity and who might be unable to keep the mouse perfectly still when clicking or double-clicking an icon. In such cases, a large threshold value may be preferred.

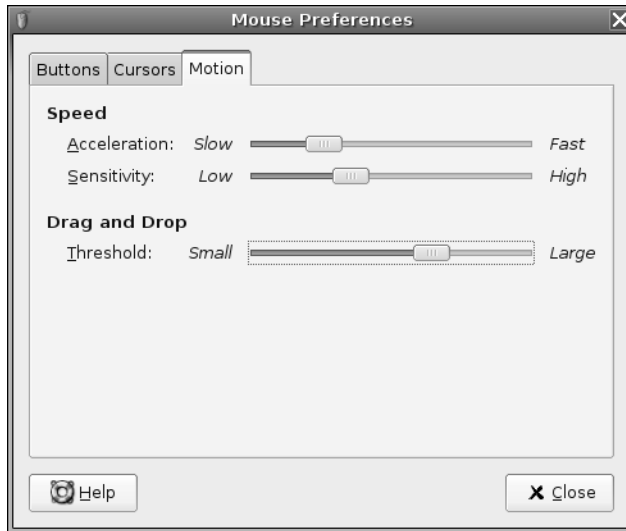


Figure 10-5. *The Mouse Preferences dialog box lets you tame that mouse.*

Changing Keyboard Settings

Select System ► Preferences ► Keyboard to open the Keyboard Preferences dialog box. This dialog box has four tabs:

Keyboard: Using the Keyboard tab, you can alter the rate of key repeat. This can be useful if you often find yourself holding down the Backspace key to delete a sentence; a shorter setting on the Delay slider and a faster setting on the Speed slider can help. However, if you get the settings wrong, you may find double characters creeping into your documents; typing an *f* may result in *ff*, for example. Beneath the Repeat Keys setting is the Cursor Blinking slider. Altering this may help if you sometimes lose the cursor in a document; a faster speed will mean that the cursor spends less time being invisible between flashes.

Layouts: On the Layouts tab, you can choose to add an alternative keyboard layout, as shown in Figure 10-6. For example, if you write in two different languages on your keyboard, it may be helpful to be able to switch between them. Click the Add button and select the second language from the list.

Layout Options: This tab lets you select from a variety of handy tweaks that affect how the keyboard works. For example, you can configure the Caps Lock key to act like a simple Shift key, or you can turn it off altogether. You can configure the Windows key so that it performs a different function, too. Put a check alongside the option you want after reading through the extensive list of options.

Typing Break: This tab features a function that can force you to stop typing after a predetermined number of minutes. It does this by blanking the screen and displaying a “Take a break!” message. Note that a notification area icon will appear before the break time to give you advanced warning of the lockout.

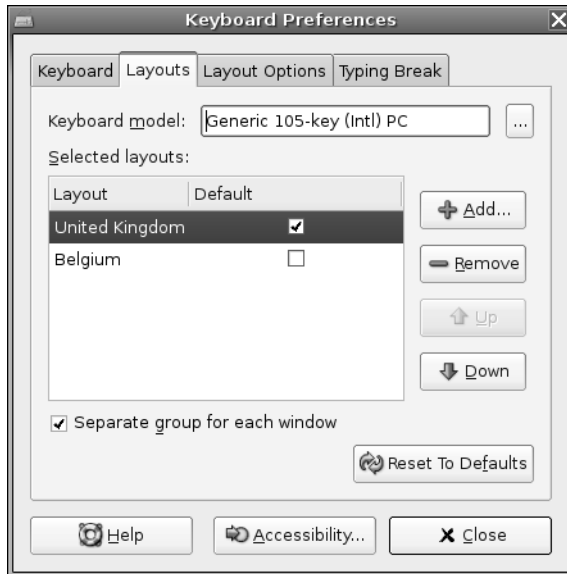


Figure 10-6. You can have more than one language setting in place for a keyboard, which is handy if you need to type in a foreign language.

Creating Keyboard Shortcuts

Ubuntu lets you define your own keyboard shortcuts for just about any action on the system. To create a shortcut, select **System** ► **Preferences** ► **Keyboard Shortcuts**. In the dialog box, search through the list for the action you want to create a shortcut for, click it, and then press the key (or key combination) you want to use. For example, you might locate the **Volume Up** and **Volume Down** entries in the list, click each, and then press **Ctrl+left arrow** and **Ctrl+right arrow**. Then you will be able to turn the volume of your sound card up or down by holding down **Ctrl** and tapping the left or right arrow key, respectively.

Caution Be careful not to assign a shortcut to a popular key. It might be nice to make Totem Media Player appear when you hit the spacebar, for example, but that will mean that it will start up several times whenever you type a sentence in a word processor! Also be aware that some key combinations are used by applications. Within OpenOffice.org's Writer, for example, the **Ctrl+left/right arrow** key combination moves you from word to word in a paragraph. If you define those combinations as shortcuts, you will no longer have this functionality.

I like to configure my `/home` folder to appear whenever I press the Home button on the keyboard. This can be done by locating the Home Folder option under the Desktop heading.

Personalizing Login Options

You can even personalize the login screen under Ubuntu. This is known technically as the GNOME Display Manager, or GDM. To access its configuration options, select System ► Administration ► Login Screen Setup. The dialog box has eight tabs:

General: In the Local drop-down list, you can choose the type of login screen: Themed Greeter, which is to say one that includes the Ubuntu graphic, or GTK+ Greeter, which is a basic login box into which you can type your details. By unchecking the Use Default Welcome boxes, you can type your own Welcome and Remote Welcome text (the remote text is what appears if someone logs in graphically via GDM across a network; it isn't something you should worry about). The Automatic Login heading lets you do away with the login box completely and go straight to the desktop from bootup. Simply put a check in the box and provide the login username. This presents obvious security issues, but if you're the only person using the computer and if it's located in a secure location, you might want to choose this option. The Timed Login option lets you select a user who will be logged in by default after a given period. This is useful if you want to present the opportunity to log in as a different user but also want to have the failsafe of logging in automatically, too.

GTK+ Greeter: Here, you can change the appearance of the simple GTK+ Greeter that will appear if you chose the relevant option on the General tab. You can change the picture by clicking the Browse button, or get rid of it completely by clicking No Logo. You can also alter the background so that there's a picture or color.

Themed Greeter: The Themed Greeter is the default under Ubuntu, and it makes the standard login page appear. You can select from a couple of other themes from the list on the left. Particularly handy is the Happy GNOME with Browser option, which shows a list of all users on your system during login. To log in, just click the appropriate entry and type the necessary password. Clicking the Install New Theme button lets you install a Themed Greeter theme from disk. You can download these from <http://art.gnome.org>.

Security: This tab lets you alter login settings that might present a security risk to your system. For example, if the Happy GNOME with Browser option is activated in the Themed Greeter tab, you can activate the Face Browser, which will show a picture on the login screen. This is considered a security risk, because it removes the need for people to type in their usernames, thus potentially handing valuable information to hackers. You can also activate the X Display Manager Control Protocol (XDMCP) for GDM, which will let users log in graphically using GDM from a remote computer (this requires the remote system to be specially configured).

XServer: This tab lets you specify settings that can be passed to the X server when it's started by GDM. In most cases, you can ignore this tab.

Accessibility: Here, you can allow activation of the GNOME Assistive Technology modules (see sidebar on following page). You can also have GDM make sounds when various events occur, such as a bad login. This can help partially sighted people.

XDMCP: If you activated XDMCP for GDM in the Security tab, this tab lets you fine-tune various settings, such as the network port it listens on as well as the maximum number of remote sessions. Unless you have specifically activated XDMCP, you can ignore this tab.

Users: Here, you can specify which users are offered as choices within GDM if the Face Browser option is activated in the Security tab. Bear in mind that Linux has many system user accounts that aren't designed to allow logins. By default, all users who have a password are displayed, which is the best way of working (the system accounts don't have passwords because they aren't login accounts).

SETTING UP ASSISTIVE TECHNOLOGY SUPPORT

You might know about the Accessibility tools under Windows, which help people with special needs use the computer. It's possible to use an on-screen magnifier so that users can better see what they're typing or reading, for example.

Under the GNOME desktop, the Accessibility tools are referred to as Assistive Technology Support. To use them, you need to install additional software packages, and then enable them. Follow these steps:

1. Assuming the Synaptic Package Manager is set up to use the online repositories (see Chapter 8), open the program (System ► Administration ► Synaptic Package Manager).
2. Click the Search button and enter `gok` as a search term. In the list of results, click the check box alongside the `gok` entry, and then click Mark for Installation.
3. Click the Search button again and search for `gnopernicus`. Again, mark its entry for installation. Then click Apply.
4. Once the packages are installed, select System ► Administration ► Preferences ► Assistive Technology Support. Click the check box alongside Enable Assistive Technologies.
5. Choose from the list the features you would like to use. They will then start automatically the next time you log in. The options work as follows:
 - The Screenreader uses a speech synthesizer to announce whatever you click on, as well as whatever you type. To alter its settings, click the Preferences button in the Gnopernicus dialog box, and then click the Speech button in the Preferences dialog box.
 - The GNOME Onscreen Keyboard (GOK) can be used by a mouse, but is most useful when an alternative input device is used, such as a touch screen. As well as presenting a virtual keyboard, it shows the options on screen as a large and easy-to-activate series of buttons. For more information, click the Help button when GOK starts.
 - The Magnifier divides the screen into two halves. The right side displays a magnified version of the left side. To learn more, click the Help button in the Gnopernicus dialog box.
 - Gnopernicus also includes support for Braille output devices. To learn more, click the Help button.

Changing Your Login Picture

If you activate the Face Browser feature along with the Happy GNOME with Browser option, GDM will display a picture alongside your username on the login screen, as shown in Figure 10-7. You can then click this and type your password to log in. You might be familiar with a similar system under Windows XP.

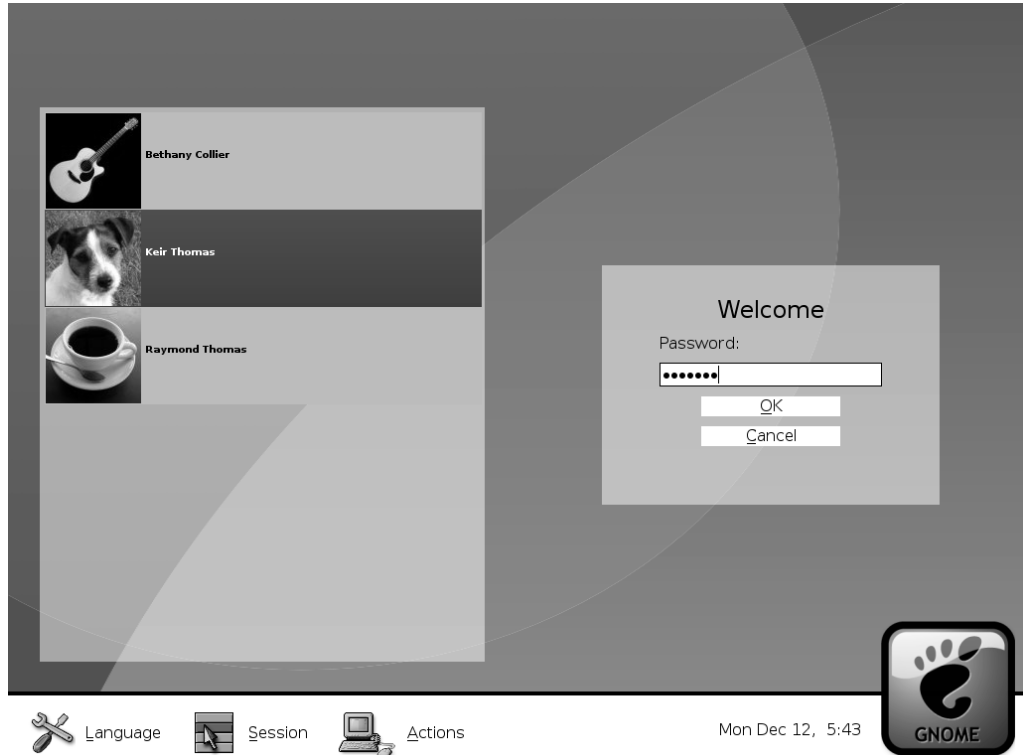


Figure 10-7. The Face Browser lets each user choose an icon to appear on the login screen.

You can choose your login picture by clicking System ► Preferences ► Login Photo. Ubuntu offers several pictures to choose from, or you can click the Browse button to locate your own. Ideally, the image you choose should be square and 96×96 pixels, although if the picture is too large, it will be automatically scaled down.

Adding and Removing Desktop Items

Virtually the entire Ubuntu desktop can be redesigned and restructured. You can move the Applications menu from the top of the screen to the bottom to be more like Windows, for example, or you can add numerous desktop shortcuts to popular applications and/or files.

Adding a Shortcut

Ubuntu's nearest equivalent to a Windows-style desktop shortcut is a *launcher*. An important difference, however, is that launchers are designed to run a certain command. Therefore, they can only point at programs (although you could create a launcher that contained a command chain required to run a particular program and file; to use The GIMP to open a picture, you might create a launcher that ran `gimp picture.jpg`, for example).

If you want to make a shortcut to a data file, such as a picture, you need to create a *link*. This is just as easy as creating a launcher.

Creating a Launcher

You can create a launcher two ways. One way is to simply click and drag an icon from one of the main menus to the desktop. This effectively copies the menu's launcher to the desktop, rather than creating a new launcher, but the effect is the same.

Note Not all menu items can be dragged and dropped to create launchers. In particular, most items on the System menu cannot be dragged and dropped.

The other way to create a launcher is to right-click the desktop and select Create Launcher. In the Create Launcher dialog box, you need to fill in only the Name and Command field; the others fields can be left blank. The Command field must contain a Linux executable program, command, or script. If you use a command-line program or script, you must check the Run in Terminal box. This will open a GNOME Terminal window automatically and run the command or script within it. The terminal window will disappear as soon as the command has finished.

To choose an icon for your launcher, click the Icon button in the Create Launcher dialog box. If you don't choose an icon, the stock GNOME one is used (the same icon as is used for unidentified and/or system files). You can select from several predefined icons or choose your own picture by clicking the Browse button, as shown in Figure 10-8.



Figure 10-8. *Creating a launcher is easy. Just fill in the Name and Command fields, and choose an icon.*

Creating a Link

Although you can create a link to a program, which will then work the same as a launcher, links are usually used to create shortcuts to files. If you then double-click the shortcut, the application associated with that file type will open the file. If you create a shortcut to a picture, for example, when you double-click it, the Eye of GNOME previewer will start.

To create a link, locate the file you want to create the link to, right-click it, and select Make Link. Then copy the new link to wherever you want it to appear, such as the desktop. You don't need to choose an icon, because the link inherits the icon of the original file. For example, if it's a picture link, it will inherit the thumbnail preview icon.

Note If you find the Make Link option grayed out, it's likely that you don't have sufficient permissions to write the link to the directory in question.

Adding and Removing Menus

You can add either the Applications menu or the entire set of menus (Applications, Places, and System) to the bottom of the screen. This can help those who long for the Windows Start button approach to access programs.

Adding All the Menus to a Panel

To add the Application, Places, and System menu to the Panel at the bottom of the Ubuntu desktop, follow these steps:

1. Right-click a blank spot on the bottom Panel and select Add to Panel.
2. In the dialog box that appears, click the Menu Bar option to add all three menus.
3. Click the Add button at the bottom of the dialog box.

Adding the Applications Menu to a Panel

When you add just the Applications menu, it appears as a single icon. To move it to the bottom-left side of the desktop, in the same location as the Windows Start menu, you need to delete and then re-create the Show Desktop icon and the Window List. Here are the steps:

1. Right-click a blank spot on the bottom Panel and select Add to Panel.
2. In the dialog box that appears, click the Main Menu option.
3. Click the Add button at the bottom of the dialog box.
4. Right-click the Window List's grab bar and select Delete.
5. Right-click the Show Desktop icon and select Delete.
6. Right-click the new menu icon, select Move, and relocate it to the bottom left.
7. Right-click the Panel and select Add to Panel.
8. In the Desktop & Windows category, select Window List, and then click Add.
9. Repeat steps 7–8 to add the Show Desktop button.

When you're finished, the Applications menu should appear as shown in Figure 10-9.



Figure 10-9. *If you just can't do without that Start button, you can re-create one on your Ubuntu desktop.*

Deleting a Menu

Creating new instances of the menus won't delete the old ones. If you create a new Applications menu at the bottom of the screen, for example, the old Applications menu will remain at the top of the screen. In fact, you can have as many instances of the menus on the desktop as you wish, although this won't be a good use of desktop space!

To delete any menu, simply right-click anywhere on that menu and select Remove from Panel.

GIVE ME MY TRASH CAN!

The developers who designed Ubuntu's desktop decided to keep the desktop clean of icons. This included relegating the Wastebasket icon to its own applet at the bottom-right side of the screen. Many people find using the applet a little difficult and miss the desktop trash can icon, which has been present on Windows and Mac OS desktops for more than 20 years.

The good news is that it's easy to get the trash can back. Click Applications ► System Tools ► Configuration Editor. In the program window that appears, click the down arrows next to Apps, then Nautilus, and then Desktop. On the right side of the program window, put a check in the `trash_icon_visible` entry.

Alternatively, in the Configuration Editor, click Edit ► Find and enter `trash_icon_visible` as a search term. Make sure that the Search Also In Key Names box has a check in it. Then click Find. The results will be listed at the bottom of the program window. Click the `/apps/nautilus/desktop/trash_icon_visible` entry. Then make sure there's a check in the `trash_icon_visible` box.

Be careful when using the Configuration Editor program. It lets you configure just about every aspect of the GNOME desktop and doesn't warn you when you're about to do something devastating, so the potential for accidental damage is high!

Personalizing the Panels

Panels are the long strips that appear at the top and bottom of the Ubuntu screen and play host to a choice of applets and/or icons. You can add a new panel by simply right-clicking an existing one and selecting New Panel, or you can remove a panel by right-clicking it and selecting Delete This Panel.

Caution If you delete a panel, the arrangement of applets it contains will be lost. Of course, you can always re-create the collection of applets on a different panel.

By right-clicking a panel and selecting Properties, you can change its size and dimensions. For example, by unchecking the Expand box, you can make the panel shrink to its smallest possible size. Then, when you add new applets (or, in the case of a panel containing the Window List, a new program is run), the panel will expand as necessary. This can be a neat effect and also creates more desktop space.

Selecting the Autohide feature will make the panel slide off the screen when there isn't a mouse over it. Choosing Show Hide Buttons will make small arrows appear on either side of the panel so that you can click to slide it off the side of the screen when it's not in use.

Working with Applets

Almost everything you see on the desktop is an applet, with the exception of shortcut icons and the panels. A menu is a form of applet, for example, as is the Workspace Switcher.

Ubuntu provides many more applets that you can choose to add to the desktop to provide a host of useful or entertaining functionality. To add an applet, right-click the Panel and select Add to Panel. As shown in Figure 10-10, you have a wide choice of applets, divided into categories. Many require configuration when they've been created, so you may need to right-click them and select Properties. For example, you'll need to set your location in the Weather Report applet's properties so it can provide accurate forecasting.

To remove an applet, simply right-click it and select Remove from Panel.

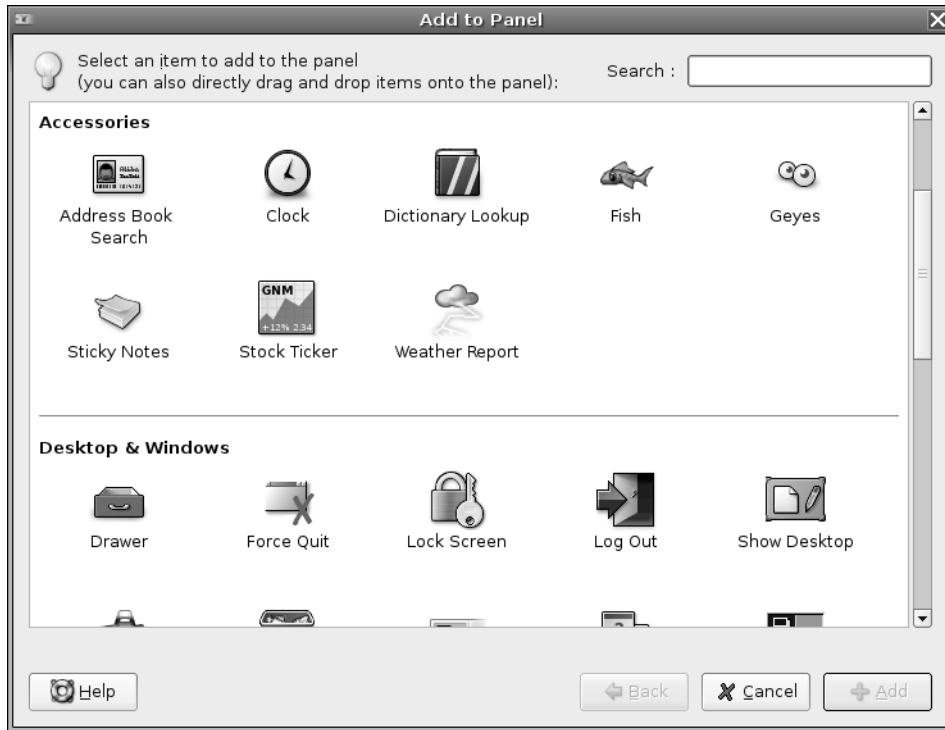


Figure 10-10. A wide variety of applets are available. Some are informative; others are just fun.

Setting Power-Saving Features

Ubuntu includes a number of features to use the power-saving features of your computer, including the ability to switch off the monitor after a set period of inactivity. However, some quick configuration is necessary to set up the system just the way you want it.

Tip If your computer has a CPU that can adjust its clock speed on the fly, such as a mobile processor or an AMD chip with the PowerNow! function, Ubuntu will automatically install software that will make this work. To see the speed of your processor, right-click the Panel, select Add to Panel, and choose the CPU Frequency Scaling Monitor.

Controlling the Monitor Shutoff

You can configure the monitor to go into standby mode after a certain amount of time has passed. This can save a lot of electricity should you happen to leave your computer unattended for long periods.

To configure this aspect of Ubuntu, select **System** ► **Preferences** ► **Screensaver**, and click the **Advanced** tab, as shown in Figure 10-11. In the **Display Power Management** section, make sure the **Power Management** box is checked. Set the desired timings in the **Standby**, **Suspend**, and **Off** boxes. Not all monitors support all three features. Consult your monitor's manual to determine which you should configure. If you're unable to find this information, simply set all three to similar values. For example, set **Standby** to 30 minutes, **Suspend** to 31, and **Off** to 32.

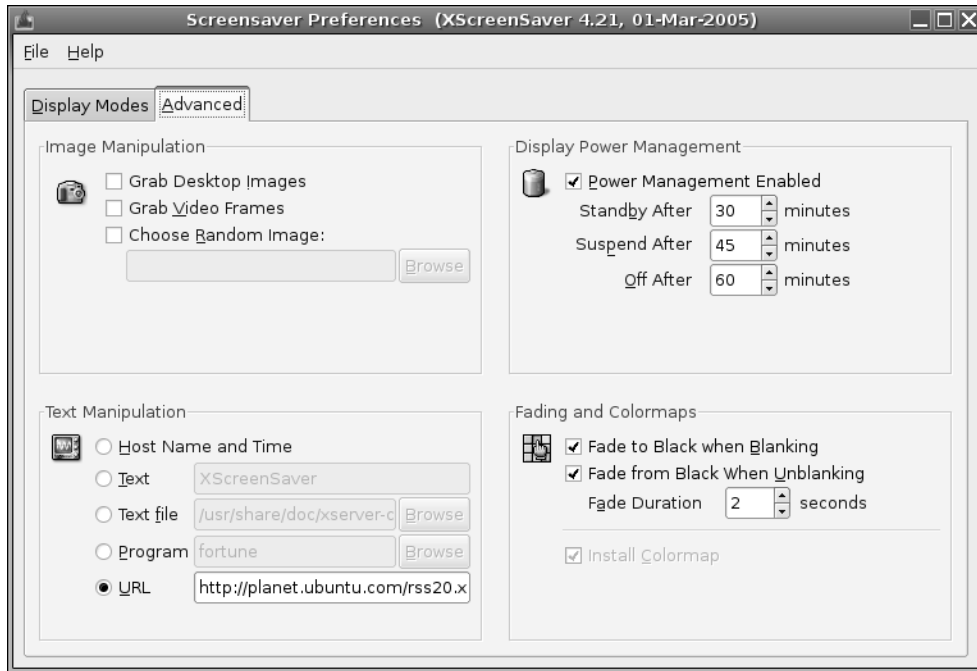


Figure 10-11. You can configure your monitor to automatically switch itself off after a certain period.

Spinning Down the Hard Disk

All modern hard disks come with the ability to spin down their motors to save energy. Then, when data is requested, the motors spin up again. There may be a slight delay while this happens, and some people dislike using disk spin-down because of this. However, on a notebook, it can lead to a substantial increase in battery life. On a desktop system, it's worth considering, because over the lifetime of a computer, it can save a lot of electricity (and therefore money!).

The spin-down settings are contained in the `/etc/hdparm.conf` file, which you'll need to edit by hand. Follow these steps to adjust the spin-down settings:

1. Open a GNOME Terminal window (**Applications** ► **Accessories** ► **Terminal**).
2. Type the following in the terminal window:


```
sudo gedit /etc/hdparm.conf
```
3. Click **Search** ► **Find and**, in the box, type `spindown_time`.

4. Click the Find button, and then close the Search dialog box.
5. Change the line the Find routine has highlighted to remove the hash mark from the beginning, so it reads like this:

```
spindown_time = 24
```

You can alter the value to anything you want. Each time unit is five seconds, so 24 equates to 120 seconds (24×5 seconds), or 2 minutes. To set a time of 20 minutes, enter 240 (240×5 seconds). If you specify a number above 240, the time units are increased to 30 minutes. In other words, a value of 241 will equate to 30 minutes, a value of 242 will equate to 60 minutes, and so on.

6. When you've finished, save the file.
7. Reboot for the settings to take effect.

POWER SAVING—IS IT WORTH IT?

An average computer draws anywhere between 100 to 500 watts of power. An average light bulb draws around 150 watts of power, so you can see that, relatively speaking, computers are low power consumers compared to many household devices. However, it's still worth considering employing power-saving techniques. You might not save yourself a lot of money, but if you switch on power saving, and your neighbor does too, and her neighbor does, then the cumulative effect will add up, and we can all contribute less towards global warming.

Try to avoid leaving your computer turned on overnight, or when you're away from it for long periods. As well as saving power, switching off your computer will avoid wear and tear on its components. Although the CPU can work 24/7 without trouble, it's cooled by a fan that's a simple mechanical device. There are other fans in your computer too, such as the graphics card fan and case fan. Each of these will eventually wear out. If your graphics card fan stops working, the card itself will overheat and might burn out. The same is true of the CPU fan. However, by shutting down your computer overnight, you can effectively double the life of the fans and radically reduce the risk of catastrophic failure. Isn't that worth considering?

Summary

In this chapter, you've learned how to completely personalize Ubuntu to your own tastes. We've looked at changing the theme so that the desktop has a new appearance, and we've examined how to make the input devices behave exactly as you would like.

In addition, you've learned how to add and remove applets from the desktop in order to add functionality or simply make Ubuntu work the way you would like.

Finally, we looked at the power-saving functions under Ubuntu and how you can avoid your computer wasting energy.

In the next chapter, we will look at what programs are available under Ubuntu to replace those Windows favorites you might miss.



Ubuntu Replacements for Windows Programs

Ubuntu is a thoroughly modern operating system and, as such, includes a comprehensive selection of software for just about every day-to-day task. Regardless of whether you want to write letters, edit images, or listen to music, Ubuntu offers something for you.

This chapter introduces the software under Ubuntu that performs the tasks you might be used to under Windows. It's not a detailed guide to each piece of software. Instead, this chapter aims to get you up and running with the Ubuntu replacement as quickly as possible. The chapter will tell you the name of the software, where you can find it on Ubuntu's menus, and a few basic facts about how to use it. In many cases, these applications are covered in far more depth later in the book.

Available Software

Table 11-1 lists various popular Windows programs alongside their Ubuntu counterparts. You'll find most of the programs listed on the Applications menu. Table 11-1 also includes a number of other mainstream alternatives, most of which aren't installed by default under Ubuntu, but are available from the Ubuntu online software repositories. You might want to try these later on. As you might expect, they're all free of charge, so you have nothing to lose.

Note Table 11-1 lists only a fraction of the programs available under Linux. There are quite literally thousands of others. The programs listed here are those that work like their Windows equivalents and therefore provide an easy transition.

Table 11-1. *Linux Alternatives to Windows Software*

Type of Program	Windows	Ubuntu	Alternative Choices
Word processor	Microsoft Word	OpenOffice.org Writer	AbiWord (www.abisource.com), KOffice KWord (www.koffice.org/kword)
Spreadsheet	Microsoft Excel	OpenOffice.org Calc	Gnumeric (www.gnome.org/projects/gnumeric/), KOffice KSpread (www.koffice.org/kspread)
Presentations	Microsoft PowerPoint	OpenOffice.org Impress	KOffice KPresenter (www.koffice.org/kpresenter)
Drawing (vector art)	Adobe Illustrator	OpenOffice.org Draw	Inkscape (www.inkscape.org), KOffice Karbon14 (www.koffice.org/karbon)
Database	Microsoft Access	OpenOffice.org Base	Rekall (www.thekompany.com/products/rekall/)
Web page creation	Microsoft FrontPage	OpenOffice.org Writer	Mozilla Composer (www.mozilla.com), Amaya (www.w3.org/Amaya/)
E-mail	Microsoft Outlook	Evolution	Mozilla Thunderbird KMail (http://kmail.kde.org)
Contacts manager/ calendar	Microsoft Outlook	Evolution	Kontakt (www.kontakt.org)
Web browser	Microsoft Internet Explorer	Mozilla Firefox	Konqueror (www.konqueror.org), Opera (www.opera.com) ¹
CD/DVD burning	Nero	Nautilus ²	K3B (www.k3b.org), X-CD-Roast (www.xcdrast.org)
MP3 player	Winamp	Rhythmbox	XMMS (www.xmms.org), Banshee (http://banshee-project.org)
CD player	Windows Media Player	CD Player	XMMS (www.xmms.org), AlsaPlayer (www.alsaplayer.org)
Movie/DVD player	Windows Media Player	Totem Media Player	MPlayer (www.mplayerhq.hu/homepage/)
Image editor	Adobe Photoshop	The Gimp	KOffice Krita (www.koffice.org/krita)
Zip files	WinZip	Archive Manager	TkZip (www.woodsway.com/TkZip/)
MS-DOS prompt	cmd.exe/command.exe	GNOME Terminal	Xterm (www.x.org) ³
Calculator	Calc	GCalcTool	Too many to mention!

Table 11-1. *Linux Alternatives to Windows Software*

Type of Program	Windows	Ubuntu	Alternative Choices
Text editor/viewer	Notepad	Gedit	KWrite (www.kde-apps.org/content/show.php?content=9901)
Desktop games	Minesweeper/Solitaire	Mines/AisleRiot Solitaire	Too many to mention!

¹ *Opera is a proprietary project, rather than open source; however, it is free of charge.*

² *Nautilus is the file manager within Ubuntu; to activate its CD/DVD burning mode, click Go ► CD/DVD Creator.*

³ *Xterm is part of the X.org package, so it is installed by default under Ubuntu. To use it, type `xterm` in a GNOME Terminal window. See Chapter 10 to learn how to create a permanent desktop launcher for Xterm.*

LINUX HAS IT ALL

The Ubuntu software archives contain thousands of programs to cover just about every task you might wish to do on your computer. Diversity is vitally important within the Linux world. For example, rather than offering just one e-mail program, you'll find many available. They compete with each other in a gentle way, and it's up to you which one you settle down with and use.

Part of the fun of using Linux is exploring what's available. Of course, the added bonus is that virtually all this software is free of charge, so you can simply download, install, and play around. If you don't like a program, just remove it from your system. However, don't forget to revisit the program's home page after a few months; chances are the program will have been expanded and improved in that short period, and it might be better at meeting your needs.

A Quick Start with Common Linux Programs

The remainder of this chapter outlines a handful of the programs listed in Table 11-1. My goal is to give you a head start in using each program, pointing out where most of the main functions can be found. You'll find more details about the The GIMP image editor, multimedia tools, and office applications in Parts 5 and 6 of this book.

Keep in mind that Linux doesn't aim to be an exact clone of Windows. Some of the programs will work in a similar way to Windows software, but that's not true of all of them. Because of this, it's very easy to get frustrated early on when programs don't seem to work quite how you want or respond in strange ways. Some programs might hide functions in what seem like illogical places compared with their Windows counterparts. Some patience is required, but it will eventually pay off as you get used to Linux.

Word Processing: OpenOffice.org Writer

OpenOffice.org is an entire office suite for Linux that was built from the ground up to compete with Microsoft Office. Because of this, you'll find much of the functionality of Microsoft Office is replicated in OpenOffice.org, and the look and feel are also similar. The major difference is that OpenOffice.org is open source and therefore free of charge.

OpenOffice.org Writer (Applications ► Office ► OpenOffice.org2 Writer), shown in Figure 11-1, is the word processor component. As with Microsoft Word, it's fully WYSIWYG (What You See Is What You Get), so you can quickly format text and paragraphs. This means the program can be used for elementary desktop publishing, and pictures can be easily inserted (using the Insert menu).

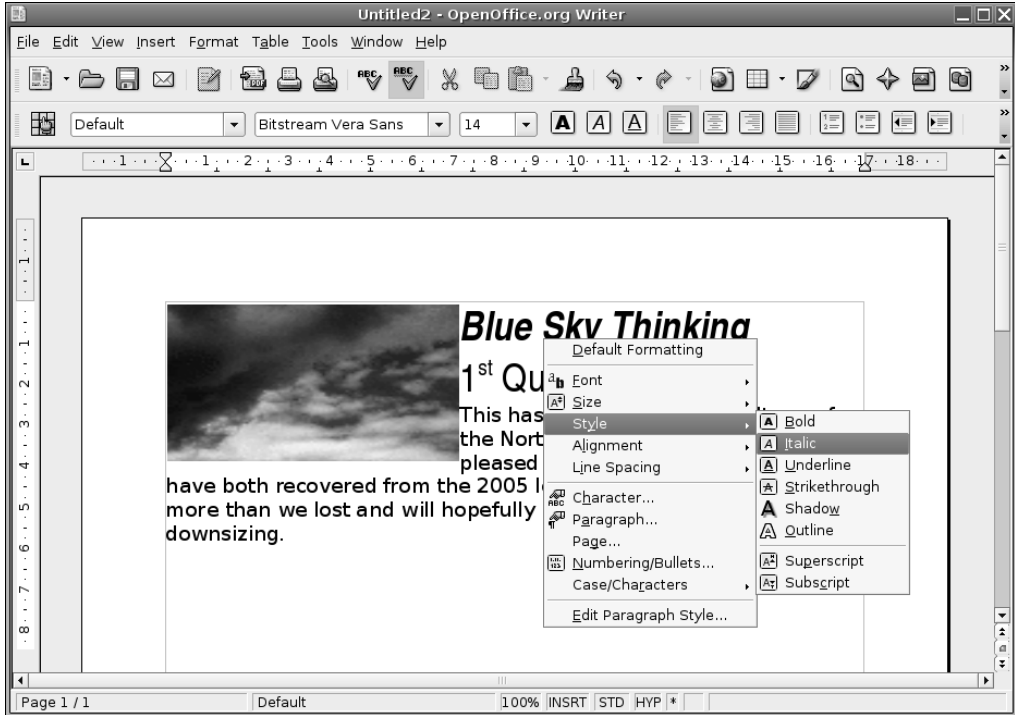


Figure 11-1. OpenOffice.org Writer

Writer's toolbars provide quick access to the formatting tools, as well as to other common functions. The vast majority of menu options match those found in Word. Right-clicking the text itself also offers quick access to text-formatting tools.

A number of higher-level functions are provided, such as mail merge and spell-checking, (found on the Tools menu). You can perform spell-checking on the fly, with incorrect words underlined in red as you type.

As with all OpenOffice.org packages, Writer is fully compatible with Microsoft Office files, so you can save and open .doc files. Just select the file type in the Save As dialog box. The only exception is password-protected Word files, which cannot be opened. You can also export documents as PDF files (using File ► Export As PDF), so they can be read on any computer that has Adobe Acrobat Reader installed.

OpenOffice.org Writer is covered in more detail in Chapter 23.

Spreadsheet: OpenOffice.org Calc

As with most of the packages that form the OpenOffice.org suite, Calc (Applications ► Office ► OpenOffice.org2 Calc) does a good impersonation of its Windows counterpart, Microsoft Excel, both in terms of powerful features and also the look and feel, as you can see in Figure 11-2. However, it doesn't run Excel Visual Basic for Applications (VBA) macros. Instead, Calc (and all OpenOffice.org programs) uses its own macro language called OpenOffice.org Basic (for more information, see <http://development.openoffice.org>).

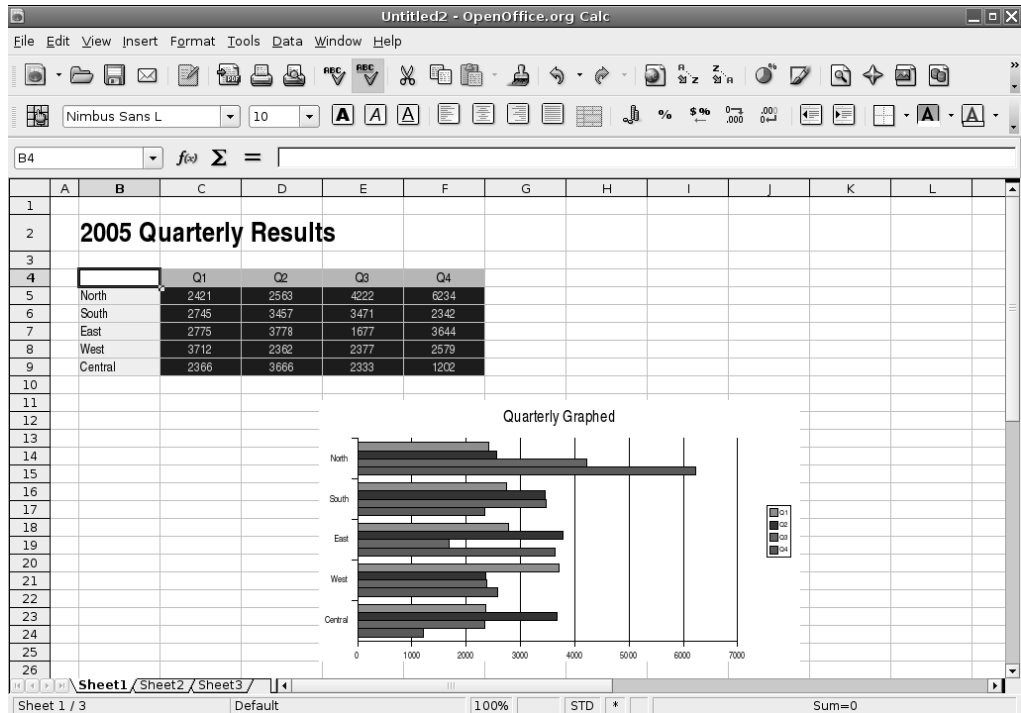


Figure 11-2. *OpenOffice.org Calc*

Calc has a vast number of mathematical functions. To see a list, choose Insert ► Function. The list includes a brief explanation of each function to help you get started. Just as with Excel, you can access the functions via the toolbar (by clicking the Function Wizard button), or you can enter them directly into cells by typing an equal sign and then the formula code. Calc is intelligent enough to realize when formula cells have been moved and to recalculate accordingly. It will even attempt to calculate formulas automatically and can work out what you mean if you type something like `sales + expenses` as a formula.

As you would expect, Calc also provides automated charting and graphing tools (under Insert ► Chart). In Figure 11-2, you can see an example of a simple chart created automatically by the charting tool.

You can format cells using the main toolbar buttons, or automatically apply user-defined styles (choose Format ► Styles and Formatting).

Tip In all the OpenOffice.org applications, you can hover the mouse cursor over each button for one second to see a tooltip showing what it does.

If you're a business user, you'll be pleased to hear that you can import databases to perform serious number-crunching. Use Insert ► Link to External Data to get the data, and then employ the tools on the Data and Tools menu to manipulate it.

As with all OpenOffice.org programs, compatibility with its Microsoft counterpart—Excel files in this case—is guaranteed. You can also open other common data file formats, such as comma-separated values (CSV) and Lotus 1-2-3 files.

OpenOffice.org Calc is covered in more detail in Chapter 24.

Presentations: OpenOffice.org Impress

Anyone who has used PowerPoint will immediately feel at home with Impress, OpenOffice.org's presentation package (Applications ► Office ► OpenOffice.org2 Impress), shown in Figure 11-3. Impress duplicates most of the common features found in PowerPoint, with a helping of OpenOffice.org-specific extras.

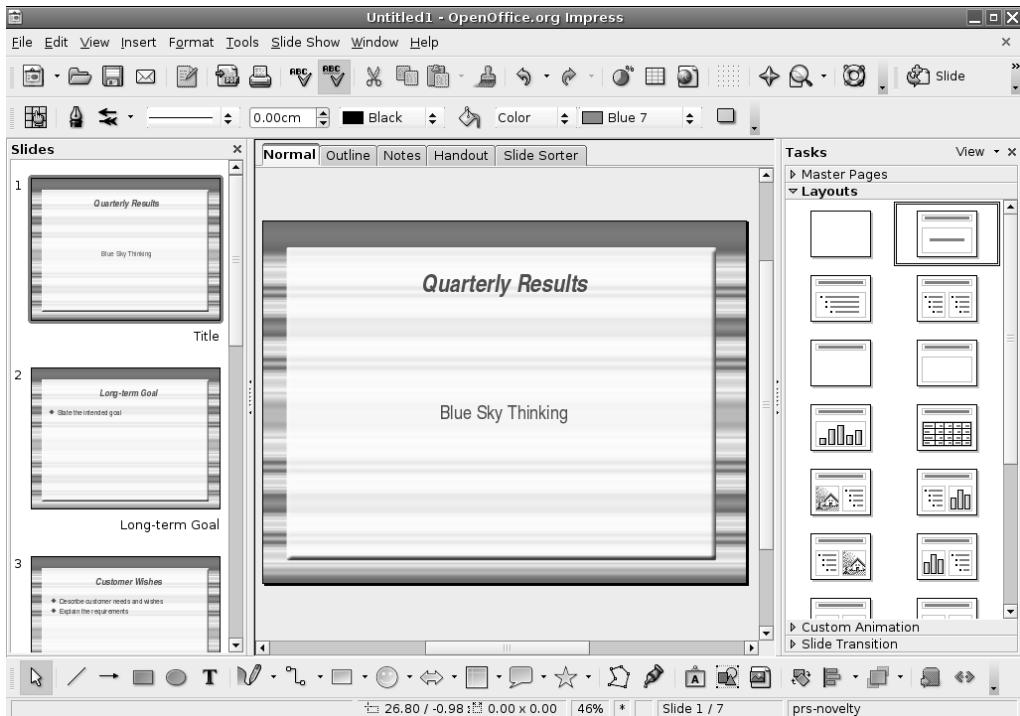


Figure 11-3. *OpenOffice.org Impress*

The program works via templates into which you enter your data. Starting the program causes the Presentation Wizard to appear. This wizard guides you through selecting a style of presentation fitting the job you have in mind. At this point, you can even select the type of transition effects you want between the various slides.

Once the wizard has finished, you can choose from the usual master and outline view modes (look under the View menu). Outline mode lets you enter your thoughts quickly, while master mode lets you type straight onto presentation slides.

You can format text by highlighting it and right-clicking it, by using the Text Formatting toolbar that appears whenever you click inside a text box, or by selecting an entry on the Format menu. Impress also features a healthy selection of drawing tools, so you can create even quite complex diagrams. These are available on the Drawing toolbar along the bottom of the screen. You can also easily insert pictures, other graphics, and sound effects.

You can open and edit existing PowerPoint (PPT) files and, as with all OpenOffice.org packages, save your presentation as a PDF file. Unique to Impress is the ability to export your presentation as a Macromedia Flash file (SWF). This means that anyone with a browser and Macromedia's Flash plug-in can view the file, either after it's put online or via e-mail. Simply click File ► Export, and then choose Macromedia Flash from the File Format drop-down list.

Along with slide presentations, Impress also lets you produce handouts to support your work.

OpenOffice.org Impress is covered in more detail in Chapter 25.

Database: OpenOffice.org Base

Base, shown in Figure 11-4, is the newest component of OpenOffice.org, introduced with version 2. Base allows you to create relational databases using a built-in database engine, although it also can interface with external databases. It's very similar to Microsoft Access in look and feel, although it lacks some of Access's high-end functions. For most database uses, it should prove perfectly adequate.

If you know the fundamentals of database technology, you shouldn't have any trouble getting started with Base immediately. This is made even easier than you might expect because, when the program starts, a wizard guides you through the creation of a simple database.

As with Access, Base is designed on the principles of tables of data, forms by which the data is input or accessed, and queries and reports by which the data can be examined and outputted. Once again, wizards are available to walk you through the creation of each of these, or you can dive straight in and edit each by hand by selecting the relevant option.

Each field in the table can be of various types, including several different integer and text types, as well as binary and Boolean values.

Forms can contain a variety of controls, ranging from simple text boxes to radio buttons and scrolling lists, all of which can make data entry easier. Reports can feature a variety of text formatting and can also rely on queries to manipulate the data. The queries themselves can feature a variety of functions and filters in order to sort data down to the finest detail.

You'll learn more about Base in Chapter 26.

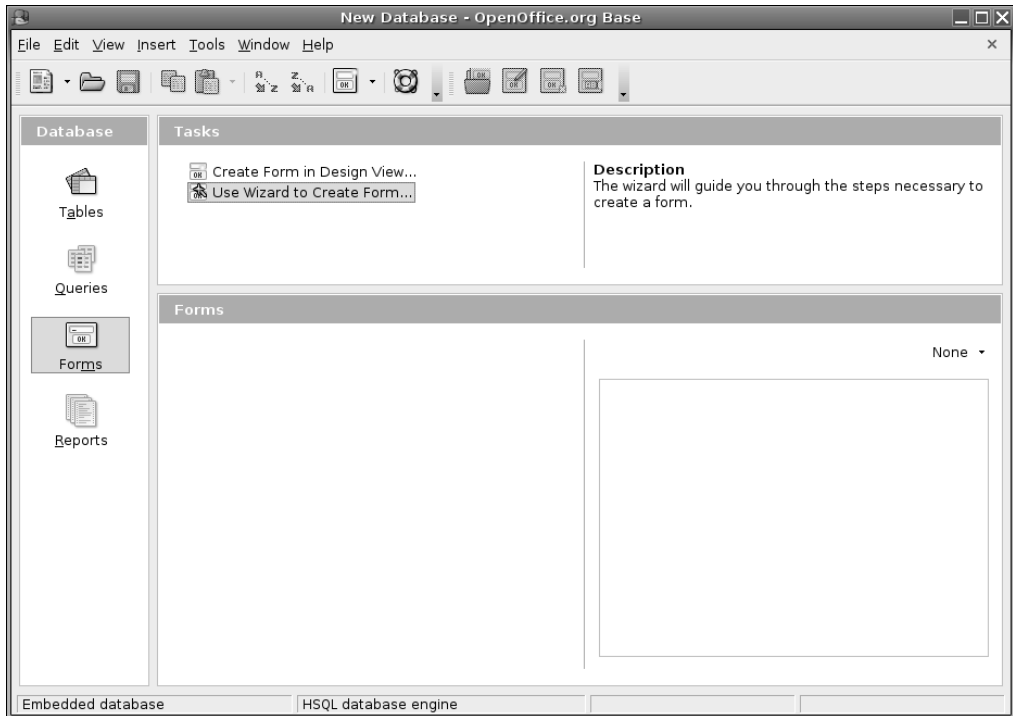


Figure 11-4. *OpenOffice.org Base*

E-Mail/Personal Information Manager: Evolution

Evolution is a little like Microsoft Outlook in that, in addition to being an e-mail client, it can also keep track of your appointments and contacts.

Before using the program, you'll need to set it up with your incoming and (if applicable) outgoing mail server settings, as detailed in Chapter 8. Evolution is compatible with POP/SMTP, IMAP, Novell GroupWise, and Microsoft Exchange servers.

Once the program is up and running, as shown in Figure 11-5, you can create a new message by clicking the New button on the toolbar. To reply to any e-mail, simply select it in the list, and then click the Reply or Reply To All button, depending on whether you want to reply to the sender or to all the recipients of the message.

To switch to Contacts view, click the relevant button on the bottom left. If you reply to anyone via e-mail, they're automatically added to this Contacts list. You can also add entries manually by either right-clicking someone's address in an open e-mail or right-clicking in a blank space in the Contacts view.

Clicking the Calendar view shows a day-and-month diary. To add an appointment, simply select the day, and then double-click the time you want the appointment to start. You can opt to set an alarm when creating the appointment, so that you're reminded of it when it's scheduled.

Finally, by clicking the Tasks button, you can create a to do list. To add a task, click the bar at the top of the list. Once an entry has been created, you can put a check in its box to mark it as completed. Completed tasks are marked with strike-through, so you can see at a glance what you still need to do.

In addition to the setup guide in Chapter 8, you'll find a full explanation of Evolution's features in Chapter 27.

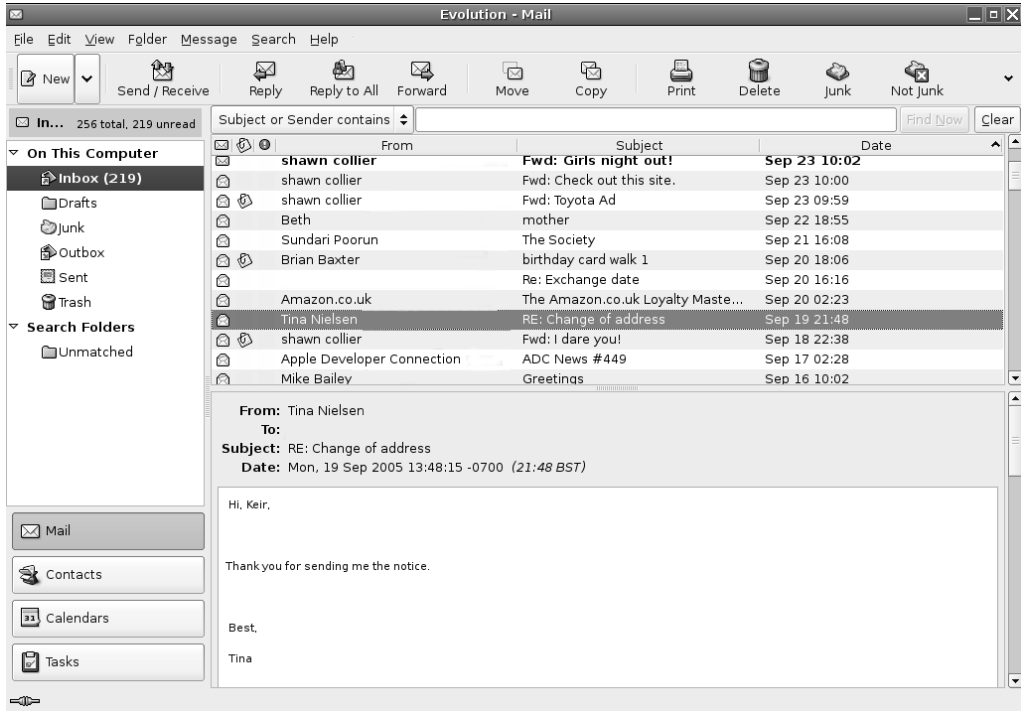


Figure 11-5. Evolution

Web Browser: Firefox

You might already know of Mozilla Firefox under Windows, where it's firmly established itself as the alternative browser of choice. The good news is that the Linux version of Firefox is nearly identical to its Windows counterpart.

When the program starts, as shown in Figure 11-6, you can type an address into the URL bar to visit a web site. If you wish to add a site to your bookmarks list, click Bookmarks ► Bookmark This Page. Alternatively, you can press Ctrl+D.



Figure 11-6. Mozilla Firefox

Searching is very easy within Firefox, using its search bar at the top right of the window. By default, Firefox uses Google for searches. To choose from other search engines, click the small down arrow on the left side of the search box. You can even enter your own choice of site if your favorite isn't already in the list.

The main benefit of Firefox over Internet Explorer is the principle of tabbed browsing, which means you can have more than one site open at once. To open a new tab, type Ctrl+T. You can move between the tabs by clicking each.

Tip When Firefox starts, tabs aren't activated. If you would like to keep tabs in view all the time, click **Edit** ► **Preferences**, and then click the **Advanced** icon. In the **Tabbed Browsing** section of the window, remove the check from "Hide the tab bar when only one web site is open."

Firefox is compatible with the same range of extensions you might have used under the Windows version of the browser. You can download new extensions from <https://addons.mozilla.org/extensions>. In addition, Firefox under Ubuntu can work with Flash animations, although you'll need to download the Flash Player software first. See the instructions in Chapter 19 to learn more.

Audio Playback: Rhythmbox and CD Player

Ubuntu's multimedia software is basic but effective. It can play back the majority of audio files, as long as it's properly configured, which is to say after additional software has been installed. I'll describe how to set up this software in Chapter 18, and if you're thinking of playing audio files on your computer, you may want to read that chapter immediately.

Rhythmbox is the audio file player. The simply titled CD Player application will automatically start whenever an audio CD is inserted into the drive. Figure 11-7 shows both of these applications.

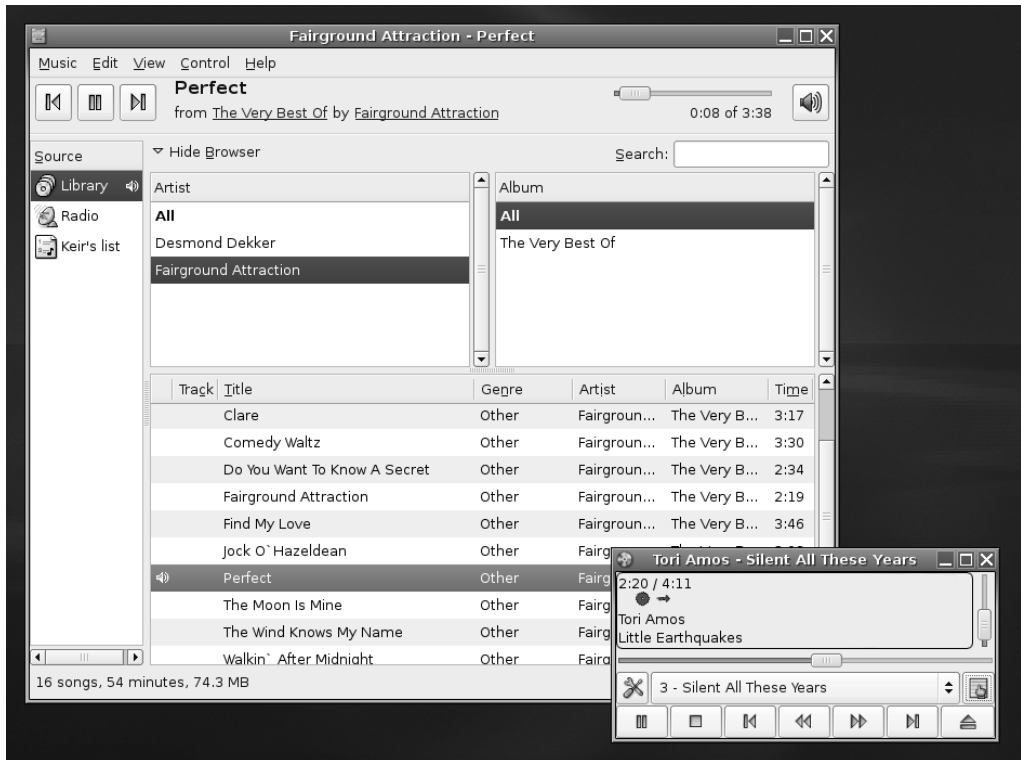


Figure 11-7. Rhythmbox and CD Player

When Rhythmbox is run for the first time, it will attempt to find and then catalog your music collection. You might be used to this kind of functionality with Windows utilities like iTunes. After the initial file search has taken place, whenever Rhythmbox runs, you will find your tracks listed by artist or name, providing they have the relevant tag information embedded in them (such as ID3 tags in MP3 music).

Note Unlike iTunes, Rhythmbox can't play Digital Rights Management (DRM)-protected files.

To start playing a music track, double-click it in the list. To make the player smaller so that it doesn't dominate the screen, click View ► Small Display.

The CD Player works just like virtually every other CD Player application you might have used. It features tape-recorder-like controls to stop and pause playback, and skip ahead/back buttons to change tracks. It is able to automatically look up the artist and track information about most CDs online, and then save the information for future reference.

You can control the output volume within the applications themselves or use the volume control applet, which is located at the top-right side of the Ubuntu desktop, near the clock. Simply click and then drag the slider to alter the volume.

Movie Playback: Totem Movie Player

Totem Movie Player is able to handle the majority of video files you might own, as long as some additional software is installed. Totem can also play back DVD movies, which, again, requires the installation of software. I'll cover setting up this software in Chapters 18 and 19; if you intend to play back videos and DVDs, these chapters should be your first port of call.

As with Rhythmbox and CD Player, Totem is a simple and uncomplicated application, as shown in Figure 11-8. The video file will play in the top left of the window. A playlist detailing movies you have queued appears in the top-right area of the program window. You can remove this, to give the video window more room, by clicking the Sidebar button.

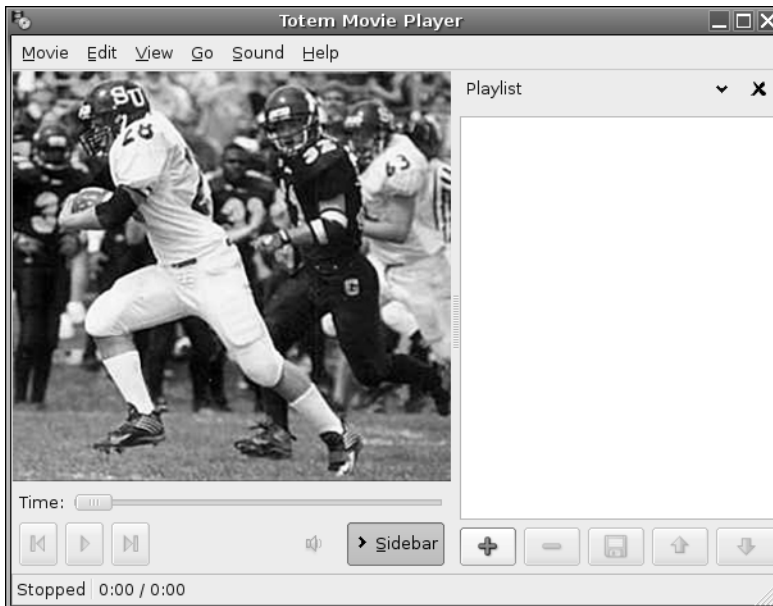


Figure 11-8. *Totem Movie Player*

You can control video playback using the tape-recorder-like controls at the bottom left. In addition, provided a compatible video format is being played, you can use the Time bar to move backward and forward within the video file. You can switch to full-screen playback by clicking View ► Fullscreen. To switch back, simply press the Escape key (labeled Esc on some keyboards). If you're watching a program that has been ripped from TV, you might want to use the Deinterlace feature on the View menu to remove any interference patterns.

Provided the software described in Chapter 19 is installed, DVD playback will start automatically as soon as a disc is inserted, and you should be able to use the mouse with any on-screen menus. In addition, you can skip between chapters on the disc using the Go menu, and also return to the DVD's main or submenu systems. To switch between the various languages on a DVD (if applicable), click Sound ► Languages and choose from the list.

CD/DVD Burning: Nautilus and Serpentine

As soon as you insert a blank writable disc, whether it's a CD or DVD, Ubuntu will detect it and offer you a range of choices. You can then create a data, music, or photo disc. If you choose to create a music CD, the Serpentine Audio-CD Creator application will open. If you choose to create a data or photo CD/DVD, a Nautilus file browser window will open in CD/DVD Creator mode.

To use Serpentine to create an audio CD, simply drag-and-drop your music files onto the program window, and then click the Write to Disc button. Keep an eye on the disc graphic at the bottom left. This is like a pie chart; when the white portion is full, the disc is full. Note that you won't be able to write certain audio files, like MP3s, to CDs unless you have the relevant codecs installed. See Chapter 18 to learn more.

Using the Nautilus CD/DVD Creator, shown in Figure 11-9, is similar to using Serpentine. Just drag-and-drop files onto the window to create shortcuts to the files. When it comes time to burn, Nautilus will take the files from their original locations. When you've finished choosing files, click the Write to Disc button on the toolbar. Unfortunately, you won't see a warning if the disc's file size has been exceeded until you try to write to the disc. However, by right-clicking an empty space in the Nautilus window and selecting Properties, you can discover the total size of the files. Remember that most CDs hold 700MB, and most DVD+/-R discs hold around 4.3GB (some dual-layer discs hold twice this amount; see the DVD disc packaging for details).

Tip Most modern CD/DVD recorders utilize burn-proof technology, which helps ensure error-free disc creation. To activate this, click Applications ► System Tools ► Configuration Editor. Click Edit ► Find, and type `burnproof`. Make sure there's a check in Search Also in Key Names. In the search results at the bottom of the window, click the first result (`/apps/nautilus-cd-burner/burnproof`) and make sure there's a check in `burnproof` at the top right of the window. Then close the Configuration Editor.

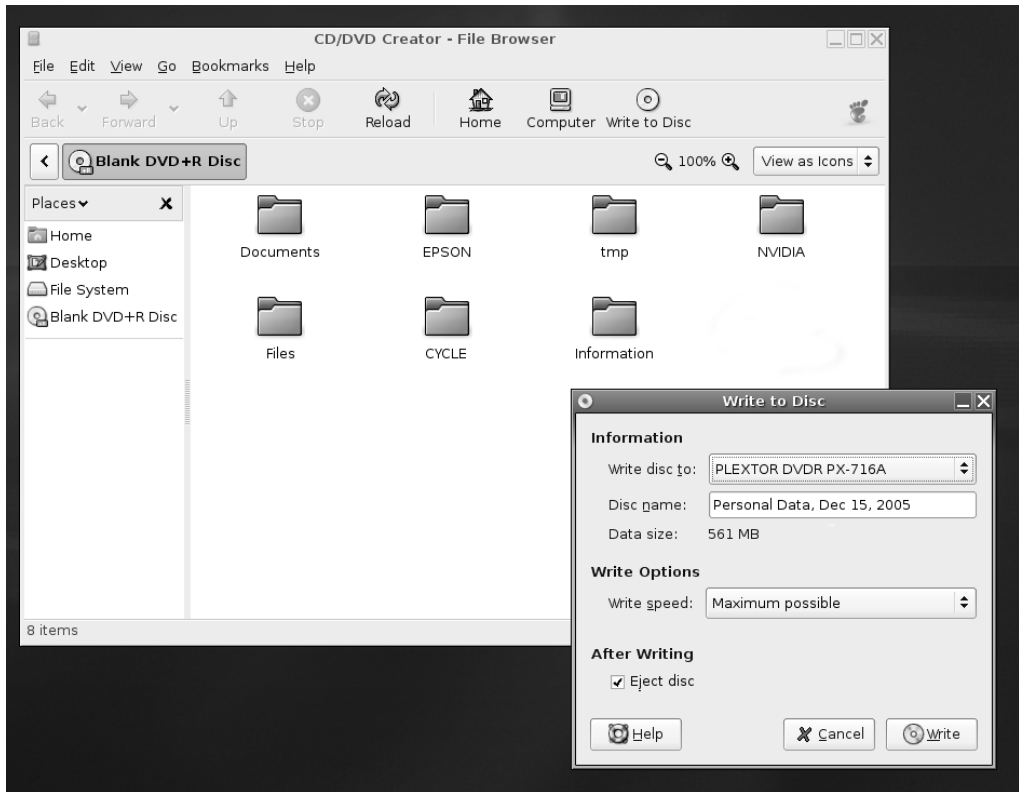


Figure 11-9. Nautilus CD/DVD Creator

Photo Editing: The GIMP

While many of the other programs introduced so far mirror the Windows look and feel in some way, The GIMP (Applications ► Graphics ► GIMP Image Editor) walks a different path. It has its own unique way of working, which takes a little getting used to. But it's very much worth the effort, because The GIMP offers photo-editing tools on par with professional products like Adobe Photoshop. It's certainly more than powerful enough for tweaking digital camera snapshots.

When you initially run The GIMP, it will install itself to your hard disk. Some of the questions it asks look complicated, but you can stick with the default choices throughout.

Once the program is running, you'll notice that it's actually little more than a large toolbar on the left side of the screen, as shown in Figure 11-10. Everything else that runs within The GIMP—whether it's a window containing the image you're editing or an additional configuration dialog box—uses its own program window. This also means that each program item that you activate gets its own button on the Panel at the bottom of the screen.




Figure 11-10. *The GIMP*

To open a picture, select **File** ► **Open** and select your image from the hard disk. Once an image file is opened, you can manipulate it using the tools on the toolbar (which are similar to those found in other image editors). On the bottom half of the main program window, you'll find the settings for each tool, which can be altered, usually via click-and-drag sliders.

To apply filters or other corrective changes, right-click anywhere on the image to bring up a context menu with a variety of options. Simple tools to improve brightness and contrast can be found on the **Layer** ► **Colors** submenu.

For an in-depth look at The GIMP package, see Chapter 20.

Other Handy Applications

Many additional applications might prove useful on a day-to-day basis. Here, I'll review some of the more common ones.

Calculator

The GNOME Calculator (also known as GCalctool) can be found on the **Applications** ► **Accessories** menu. In its default mode, shown in Figure 11-11, it shouldn't present any challenges to anyone who has ever used a real-life calculator, although the **Bksp** key might be new. This simply deletes the last number you typed (handy if you miskey during a calculation).

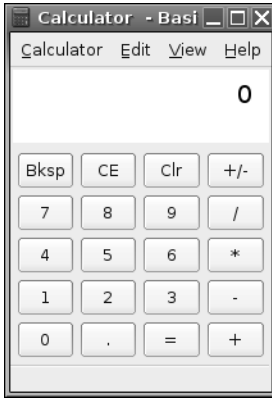


Figure 11-11. *GNOME Calculator*

Calculator also has three other modes that you can switch into using the View menu: Advanced, Financial, and Scientific. All offer calculator functions relevant to their settings. The Advanced mode is simply a more complicated version of the basic Calculator. It can store numbers in several memory locations, for example, and carry out less common calculations such as square roots and reciprocals.

Floppy Formatter

The Floppy Formatter is a small and simple program designed, as its name suggests, to format 3.5-inch floppy disks, as shown in Figure 11-12. You can find it on the Applications ► System Tools menu. The program can format disks in Linux (ext2) or DOS format. The latter is the best option, bearing in mind you might be sharing disks with others who run Windows. There's also no harm in doing so, since DOS-formatted disks will work fine under Ubuntu, too.

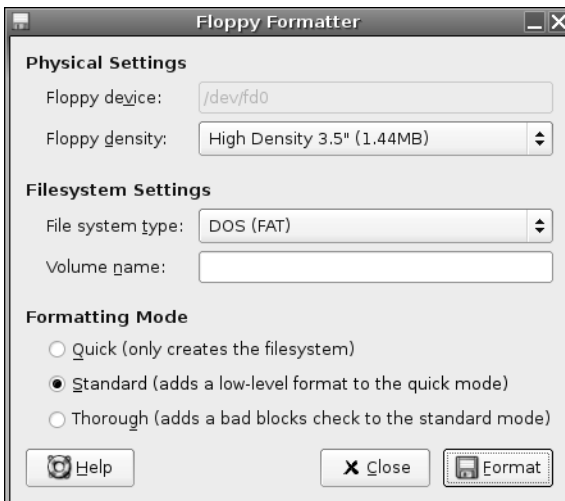


Figure 11-12. *Floppy Formatter*

You can choose between Quick, Standard, and Thorough formats. Quick simply wipes the disk's table of contents, Standard will perform an actual disk format, and Thorough adds in a disk scan to mark up any bad data blocks.

Archive Manager

Archive Manager (also known as File Roller), shown in Figure 11-13, is Ubuntu's archive tool. It's the default program that opens whenever you double-click .zip files (or .tar, .gz, or .bzip2 files, which are the native archive file formats under Linux). If you want to start the program manually, click Applications ► Accessories ► Archive Manager.

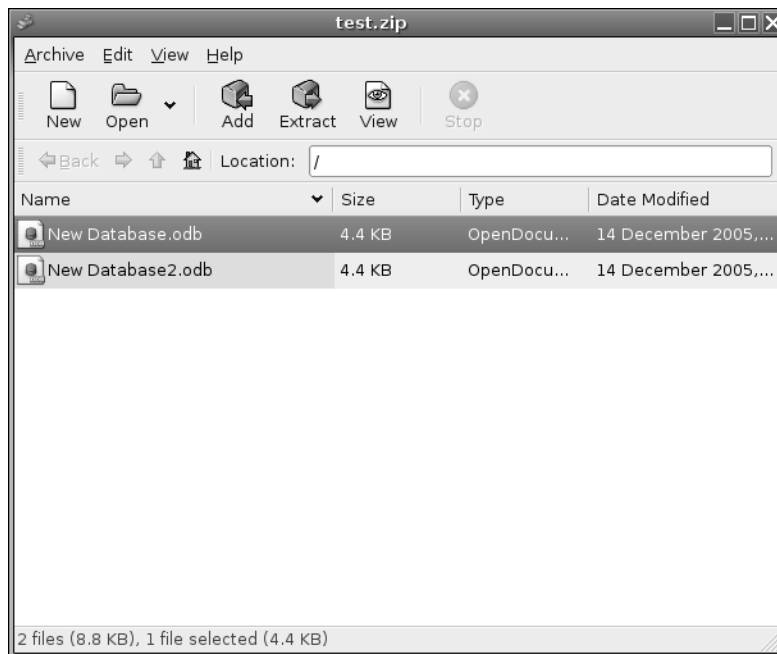


Figure 11-13. Archive Manager

To extract files from an archive, select them (hold down the Ctrl key to select more than one file), and then click the Extract button on the toolbar.

To create an archive, start Archive Manager and click the New button. Give the archive a name, and then drag-and-drop files onto the Archive Manager window. When you've finished, simply close the Archive Manager window.

Dictionary

You can use the Dictionary tool to look up the definition of words using either a dictionary stored on the hard disk or one online, at the Cambridge Dictionary web site. The dictionary on the disk is a 1913 edition of Webster's, so its word definitions are hardly current and don't take into account any recent vocabulary, but the online dictionary is certainly up-to-date.

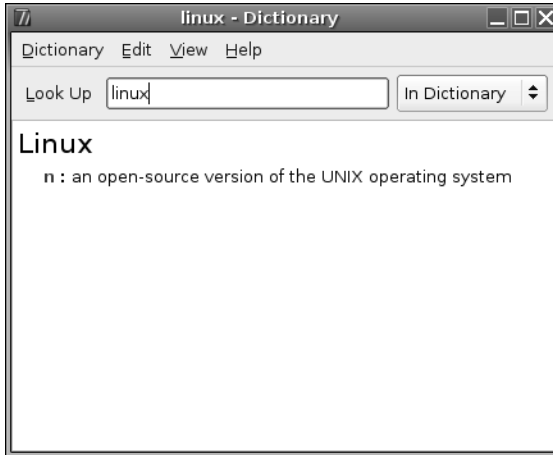


Figure 11-14. *Dictionary*

You'll find the Dictionary program, shown in Figure 11-14, on the Applications ► Accessories menu. As soon as you start typing, the program will begin to look up the word in the dictionary, and this can cause a momentary delay before the letters appear on your screen. To access the online dictionary, simply select On a Web Site from the drop-down list next to the Look Up field.

Gaim Internet Messenger

Gaim is the instant messaging software provided with Ubuntu. Unlike most other messaging programs, Gaim isn't exclusive to one chat protocol. In other words, you can use it to connect to MSN, AOL/ICQ, Yahoo!, and many other services. The program can be found on the Applications ► Internet menu.

Details for setting up Gaim are in Chapter 8. Once the program is up and running, you can chat with any of your buddies by double-clicking their icon, as shown in Figure 11-15. To set your status, click the Away icon at the bottom left and select an appropriate message.

The rest of the program can be administered by right-clicking the notification area icon that appears when the program starts. For example, you can initiate file transfers or sign off from there.

Note If you're interested in learning more about how Gaim is built, check out *Open Source Messaging Application Development: Building and Extending Gaim*, authored by Gaim project leader Sean Egan (1-59059-467-3; Apress, 2005).



Figure 11-15. *Gaim Internet Messenger*

Mines

Mines is the Ubuntu equivalent of Minesweeper and, in fact, is almost exactly the same as the Windows program, as you can see in Figure 11-16. The rules are identical, too: on each grid are several hidden mines, and it's your job to locate them. After you've clicked one square at random, you'll see a series of empty squares and several with numbers in them. Those with numbers indicate that a bomb is near. Your job is to use logic to work out where the bombs are, and then mark them by right-clicking them. Oh, and you have to do this as quickly as possible because you're being timed.

To change the grid size, click **Settings** ► **Preferences**. Your choices are Small, Medium, Large, and Custom.




Figure 11-16. *Mines*

GnomeMeeting

If you want to conduct a video conference using Ubuntu, then GnomeMeeting is for you. It supports the commonly-used H.323 protocol, as used by Microsoft's NetMeeting, and also lets you register to an Internet Locator Service (ILS) directory, as well as take part in multiple-user calls via a Multipoint Control Unit (MCU). In addition, by registering using the link provided, you can make PC telephone calls, although this requires paying a fee and setting up an account.

When the program starts, it will walk you through setup via a wizard. Simply answer the questions with the details of your setup. Once the program is up and running, as shown in Figure 11-17, type the URL of the person you would like to call into the address bar and click Call ► Connect.



Figure 11-17. *GnomeMeeting*

Summary

In this chapter, we've taken a look at some Ubuntu programs that provide vital functions that you might have used daily under Windows. The aim was to get you started with this software as quickly as possible by pointing out key features. You've seen how some programs mirror the look and feel of their Windows counterparts almost to the letter, while others resolutely strike out on their own path. It takes just a little time to become familiar with Ubuntu software, and then using these programs will become second nature.

In the next chapter, we'll move on to more fundamental Ubuntu tasks: manipulating files. However, once again, this is not too dissimilar from the Windows experience, which makes getting used to the system very easy.



Managing Your Files

Files are what make the world of Linux go round. They're the currency of any kind of operating system, because every time you use your computer, you generate new files, even if they're only temporary.

How Linux views files, as well as the disks and partitions that contain them, varies somewhat from how Windows handles files. In many ways, the Linux system of file management is far simpler than that in Windows (which, ironically, was created as an attempt to make everything easy!). The Linux system is also much more established.

In this chapter, I will explain how you can manage your files under Ubuntu. This isn't a definitive guide; you'll need to wait until Chapter 15 to learn the technical ins and outs of the file system. However, it provides enough information for you to understand how the system works, and where and how you should store your data.

Understanding File System Concepts

Just like Windows, Ubuntu has a file system that is shared among software components and your own personal data, which you generate within various applications, or perhaps download from the Internet. However, Ubuntu differs from Windows in a couple of important ways.

Drive References

Perhaps the most important differences in Linux are that it doesn't use drive letters and it uses a forward slash (/) instead of a backslash (\) in filename paths. In other words, something like `/home/john/myfile` is typical under Ubuntu, as opposed to `C:\Documents and Settings\John\myfile` under Windows. The root of the hard disk partition is usually referred to as `C:\` under Windows. In Ubuntu, it's referred to simply with a forward slash (/).

If you have more than one drive, the drives are usually combined together into the one file system under Linux. This is done by *mounting*, so that the any additional drives appear as virtual folders under the file system. In other words, you browse the other hard disks by switching to various directories within the main file system. I'll explain mounting in Chapter 15.

Case Sensitivity

Another important difference between Ubuntu and Windows is that filenames in Ubuntu are case-sensitive. This means that `MyFile` is distinctly different from `myfile`. Uppercase letters are vitally important. In Windows, filenames might appear to have uppercase letters in them but, actually, these are ignored when you rename or otherwise manipulate files.

Because of this case sensitivity, you could have two separate files existing in the same place, one called `MyFile` and another called `myfile`. In fact, you could also have `myFile`, `Myfile`, `MYFILE`, and so on, as shown in Figure 12-1.

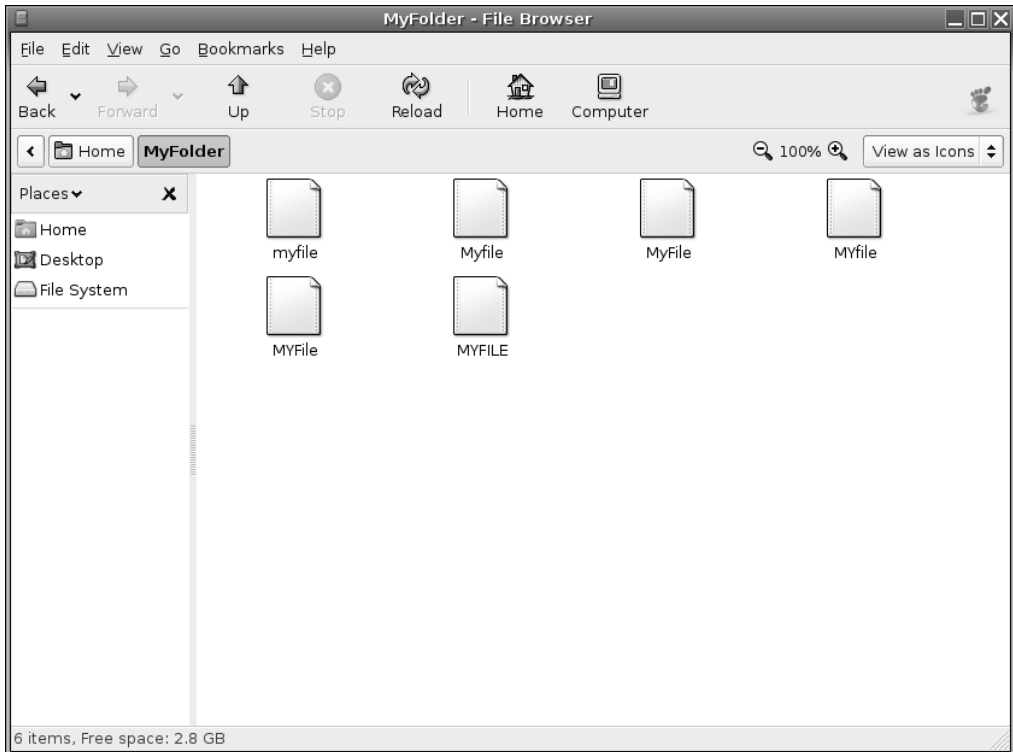


Figure 12-1. *Ubuntu filenames are case-sensitive, so many similar filenames can exist, differing only in which letters are capitalized.*

File Access and Storage

Under Windows on a desktop computer, you have access to the entire hard disk. You can write, read, or delete files anywhere (unless the system has specifically been configured otherwise). You can save your personal files in `C:\Windows`, for example. Under Ubuntu, ordinary users can browse most of the hard disk, but they aren't able to write files to the majority of folders (in some cases, they won't even be able to access files).

Although we'll cover the file system in much more depth in Chapter 15, for the moment, it's enough to know that you've been given your own part of the hard disk in which to store your stuff. This is a directory located within the `/home` directory, and its name is taken from your username. If your login name is `louisemith`, your place for storing files will be `/home/louisemith`. Figure 12-2 shows an example of a user's home directory.

Note Linux generally uses the terms *directory* and *subdirectory* for the places you put files, whereas Windows XP refers to them as *folders*. It's merely a matter of semantics. However, within the Nautilus file browser, directories are pictured as folders and are referred to as such, thus furthering the confusion!

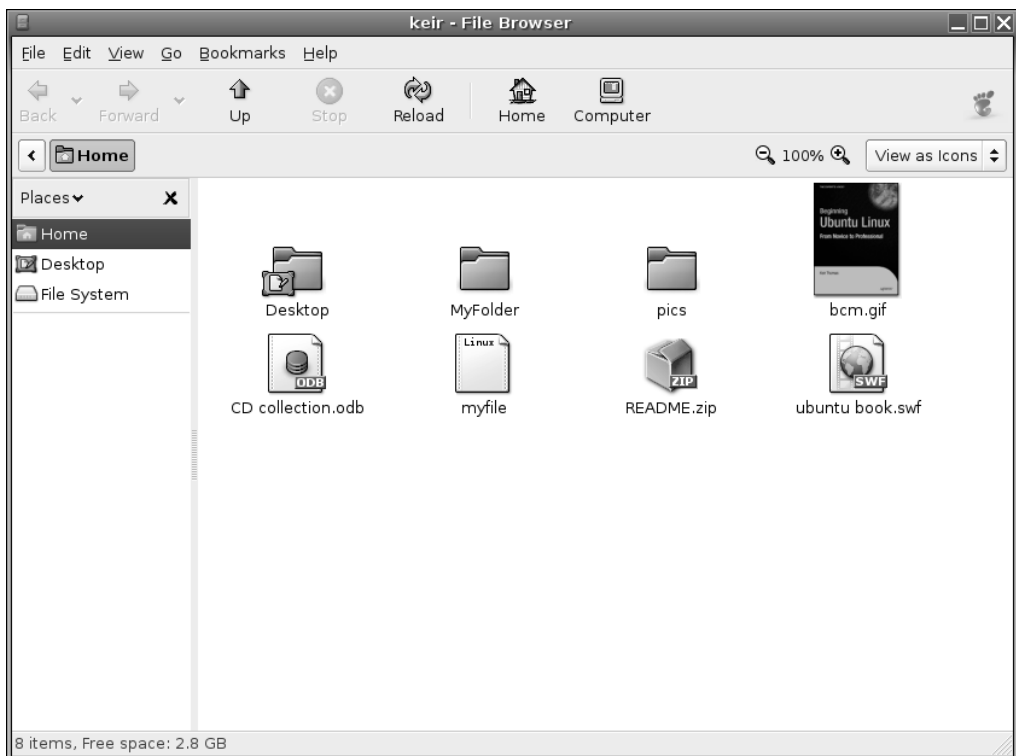


Figure 12-2. Your personal area on the hard disk is in the `/home` directory and is named after your username.

Some programs might create subdirectories in your home directory in order to store and organize their output. For example, a digital camera program might create a Pictures directory within your `/home` directory. It's up to you whether you use these. The standard practice within the Linux community is to simply save everything into your `/home` subdirectory (for example, `/home/keir`) and sort it out later!

Files within Ubuntu remember who owns them. If user `johnsmith` creates a file, he can make it so that only he can read or write the file (the default setting is that other users will be able to read the file but not write any new data to it). Directories, too, are owned by people, and the owner can set access permissions. By default, all users on a system can access each other's `/home` directories and read files, but they won't be able to change the files or write new files to any directory within `/home` that isn't theirs.

Note Any user with superuser powers has access to all of the system and can create, edit, and delete files in all directories. This is so that user can perform essential system maintenance.

Using Nautilus

Nautilus is the name of the default file browser in Ubuntu. It's not dissimilar to My Computer/Windows Explorer under Windows in that it presents a list of files on the right side of the window and a series of shortcuts to popular locations within the file system on the left side.

Starting Nautilus is simply a matter of clicking the Places menu and choosing a location, as shown in Figure 12-3. Alternatively, you can click Applications ► Accessories ► File Browser, which will open the default browsing location (your `/home` folder).

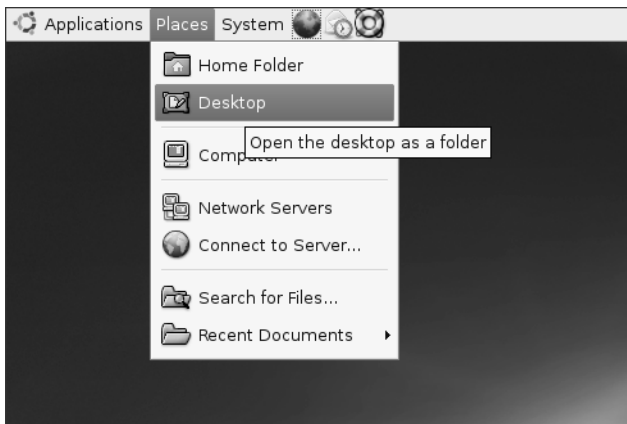


Figure 12-3. You can open a file browser window by selecting a location under the Places menu.

The Nautilus window (see Figures 12-1, 12-2, and 12-4) consists of several elements:

Menu bar: The menus offer options for controlling the way files are displayed in the Nautilus window, as well as the look and feel of Nautilus itself. The Bookmarks menu lets you create web-browser-like shortcuts to certain locations in your file system so you can access them instantly.

Toolbar: As in a web browser, the toolbar allows you to quickly move backward and forward from place to place in your browsing history. In addition, you can reload the file listing, in order to reflect any changes that might have taken place since the Nautilus window opened, and quickly navigate to popular file system locations, such as your `/home` folder.

Location bar: This feature, located beneath the toolbar, is unique to Nautilus and lets you see where you are in your file system, as well as quickly and easily move through your file-browsing history. For example, if you start in `/home/keir` and then browse to `/home/keir/pictures/holiday/disneyworld`, clicking the Pictures button will return you to `/home/keir/pictures`. The other folders listed on the location bar (`holiday` and `disneyworld` in this example) won't disappear, so you can return to those as well. It's best demonstrated by example, so give it a try!

Zoom controls: To the right of the location bar are the zoom controls. These make the icons representing the files bigger or smaller. When you're browsing a lot of files at once, shrinking them will fit more in the window. On the other hand, when you're viewing photo thumbnails, it can be handy to increase the zoom setting so you can see more detail in the pictures.

View As Icons/List: To the right of the zoom controls is a drop-down list that switches between icon and list view. List view shows details about the files, such as file size, the type of file, its permissions, and so on. Icon view presents the files as a series of large icons. In many cases, the icons will give a clue as to the nature of the file; for example, MP3 files appear with speaker and musical note graphics. If the folder you're browsing contains image files (or certain document files, such as PDFs), these will be automatically thumbnailed—the icon will be a small version of the contents of the file, as shown in Figure 12-4. This is very handy when browsing pictures for printing or editing.

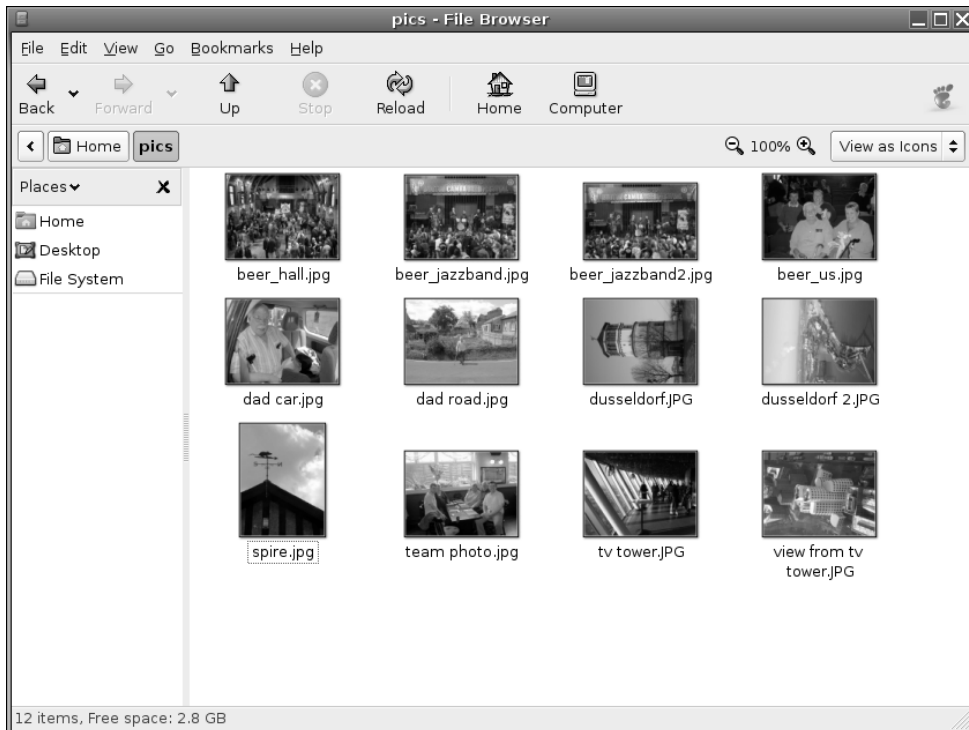


Figure 12-4. Whenever you view a folder full of pictures in icon view, they will be automatically thumbnailed.

Places pane: The Places pane on the left lists the most popular locations within the file system. Double-clicking each icon takes you to that location instantly. Clicking the File System entry takes you to the root of the file system (/).

As under Windows, you can right-click each file in the file browser window to see a context menu with options to rename the file, delete it, open it with particular applications, and so on. The Properties option on the context menu lets you view information about the file and alter certain aspects of it, such as its access permissions (discussed in Chapter 15). You can even add some text notes about the file if you wish!

Caution You should never delete your /home folder. Doing so will most likely destroy your personal Ubuntu setup and prevent you from logging in.

File and Folder Icons

You can change the icon for any file or folder to whatever you want. Right-click the file or folder, and then click the Select Custom Icon option. You can choose from a wide range of supplied icons or click the Browse button and locate your own graphic. Virtually any image can be used, regardless of format or even size, so you can use digital camera snapshots if you wish.

Files and folders can also have *emblems* applied to them. These are smaller icons that are “tagged on” to the larger icons in both list view and icon view. Emblems are designed to give you quick clues about the nature of the file. To apply an emblem, right-click the file or folder, select Properties, and then click the Emblems tab. As shown in Figure 12-5, a range of icons is available; in fact, any file or folder can have several emblems applied at once. Simply put a check in the box beside the icons you wish to apply.

Note Nautilus makes use of a handful of emblem icons for its own needs, too. For example, a file with a lock emblem attached to it indicates that you don’t have the necessary file permissions to edit or delete that file. An X emblem means you don’t have permissions to access that file or folder at all, not even to view it. In most cases, the emblems are self-explanatory.



Figure 12-5. A variety of miniature emblems can be applied to an icon to aid recognition of the file.

Special Nautilus Windows

As well as letting you view your files, Nautilus has a number of *object modes*. This is a complicated way of saying that Nautilus lets you view things other than files.

The most obvious example of this is the computer view of your file system, which presents an eagle's eye view of your storage devices. To access this view, click **Places** ► **Computer**. If you have a card reader attached, it will appear here, as will any Windows partitions that may be on your hard disk. Double-clicking each item opens a standard Nautilus file browser window (for this to work with Windows partitions, they must be set up correctly, as described in the "Accessing Windows Files" section later in this chapter).

Another Nautilus object mode is the fonts view, which lets you see at a glance any fonts installed on your computer. To access fonts view, click **Go** ► **Location** in any open Nautilus window, and then type `fonts://`.

Object mode come into its own when viewing network locations. Clicking **Places** ► **Network Servers** brings up the browsing network object view, for example. You can also browse to FTP sites by clicking **Go** ► **Location** in a file browser window and entering an FTP address (prefacing it with `ftp://`).

Note You might be used to dragging-and-dropping files onto program windows or taskbar buttons within Windows in order to open the file. This works with only some programs within Ubuntu. Generally, the best policy is to try it and see what happens. If the program starts but your file isn't opened, it obviously didn't work.

HIDDEN FILES AND DIRECTORIES

When you view your `/home` directory via Nautilus, you're not seeing every file that's there. Several hidden files and directories relating to your system configuration also exist. You can take a look at them by clicking View ► Hidden Files in the Nautilus menu. Clicking this option again will hide the files and directories.

You might notice something curious about the hidden items: they all have a period before their filenames. In fact, this is all that's needed to hide any file or directory: simply place a period at the front of the filename. There's no magic involved above and beyond this.

For example, to hide the file `partypicture.jpg`, you could simply right-click it and rename it `.partypicture.jpg`. You'll need to click the Reload button on the toolbar for the file view to be updated and for the file to disappear. As you might expect, removing the period will unhide the file.

Files are usually hidden for a reason, and it's no coincidence that most of the hidden files are system files. In addition, every program that you install, or is installed by default, will usually create its own hidden folder for its system configuration data. Deleting such files by accident can be catastrophic.

Launching Files and Running Programs

As with Windows, most of the programs on your Ubuntu system automatically associate themselves with various file types that they understand. For example, double-clicking a picture will automatically open the Eye of GNOME viewer application, and double-clicking a `.doc` file will start OpenOffice.org Writer.

Ubuntu is automatically set up to view common file types. However, you might find Table 12-1 useful. It shows which programs are required for viewing certain types of documents.

Note Whenever you install new software from the installation CD or the official software repositories, it should add an entry to the Applications menu. If for some reason this doesn't happen, you can create a shortcut using the techniques explained in Chapter 10.

Table 12-1. *Common File Types*

File Type	File Extension	Viewer	Location on Applications Menu
Word processor document	.doc, .rtf	OpenOffice.org Writer	Office ► OpenOffice.org2 Writer
Spreadsheet	.xls	OpenOffice.org Calc	Office ► OpenOffice.org2 Calc
Presentation	.ppt	OpenOffice.org Impress	Office ► OpenOffice.org2 Impress
PDF file	.pdf	Evince	Not on Applications menu ¹
Compressed file	.zip, .tar, .gz, .bz2, and others	File Roller	Accessories ► Archive Manager
Image file	.jpg, .gif, .tif, .bmp, and others	Eye of GNOME	Graphics ► Image Viewer
HTML file	.htm, .html	Firefox	Internet ► Firefox Web Browser
Text file	.txt, .log	Gedit	Accessories ► Text Editor
Audio file	.wav	Rhythmbox	Sound & Video ► Rhythmbox Music Player
Video file	.mpg, .mpeg, .avi	Totem	Sound & Video ► Totem Movie Player

¹ For some reason, Evince is not present on the Applications menu, although it is installed by default and associated with PDF files. If you wish, you can add your own shortcut following the instructions in Chapter 10.

If you want to change the program associated with a file type, right-click the file, select Open with Other Application, and choose the other program, as shown in Figure 12-6. From that point on, every time you right-click, you'll be offered the choice of the program to open the file.

Note Under Windows, you can use Windows Explorer to launch program executables by just browsing to their location within Program Files and double-clicking their .exe file. It's technically possible to run programs by browsing to their location using Nautilus, but this is discouraged. One reason is that Ubuntu doesn't store all of its programs in one central folder, as does Windows. However, most programs that are used on a daily basis can be found in /usr/bin. If the program itself isn't stored in /usr/bin, it will contain a symbolic link (effectively, a shortcut) to the program's genuine location on the hard disk.

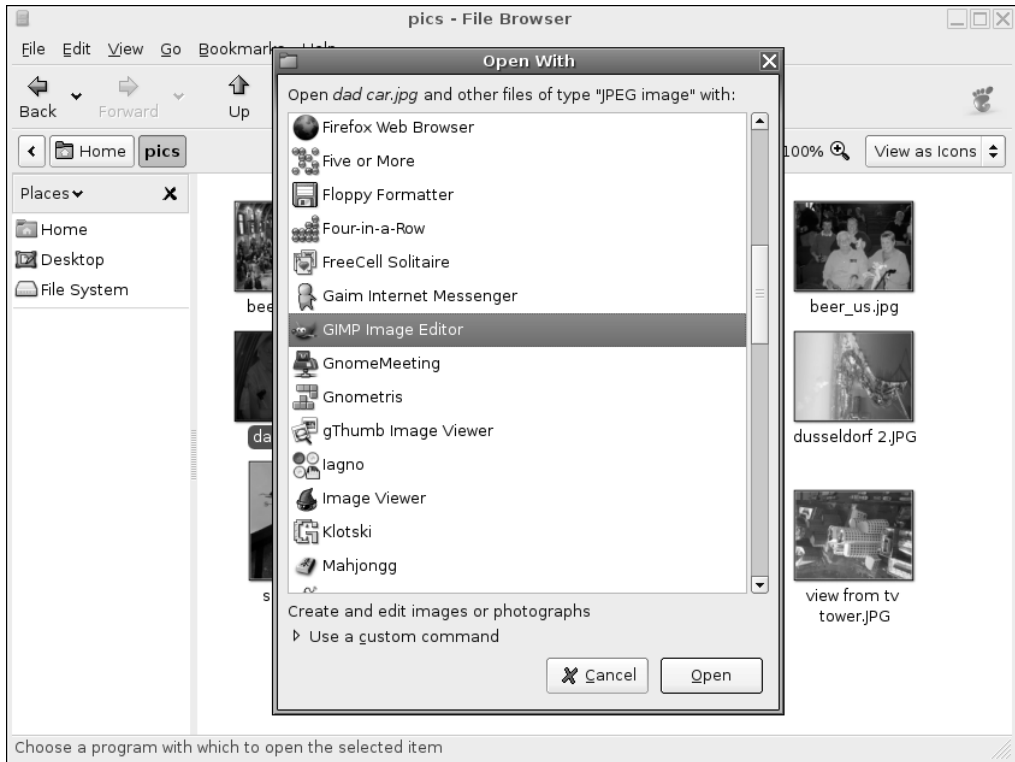


Figure 12-6. You can select which program to use to open a file by right-clicking and selecting *Open with Other Application*.

Accessing Windows Files

Running Ubuntu on your PC makes you a relative stranger in a world of Windows users. It's likely that you'll need to access Windows files on a regular basis. If you've chosen to dual-boot with Windows, you might want to grab files from the Windows partition on your own hard disk. If your PC is part of a network, you might want to access files on a Windows-based server or workstation on which a shared folder has been created.

Working with Files in Windows Partitions

Ubuntu may or may not make your Windows partitions available automatically. It depends on whether your Windows partition is FAT32, as used with Windows 95, 98, and Me, or NTFS, as used with Windows NT, 2000, and XP. FAT32 partitions will be instantly accessible, but making NTFS partitions accessible requires some work.

Note It's possible for an installation of Windows 2000 or XP to use FAT32 instead of NTFS, but this requires the user to make a deliberate choice during setup. Unless you know your Windows 2000 or XP system has been formatted with FAT32, it's very likely that it is NTFS.

If the drive is available, an icon for it should appear on the desktop. Unfortunately, with Ubuntu 5.10, NTFS drives appear on the desktop, but when you try to use them, they aren't accessible because of incorrect permissions. You can get around this bug in a couple ways. The easiest is to use the Disks Manager program, as follows:

1. Start the Disks Manager by selecting System ► Administration ► Disks.
2. Select Hard Disk in the list on the left and click the Partitions tab on the right. Then click the partition relating to your Windows installation.
3. You need to mount the partition. I discuss mounting in detail in Chapter 15, but for the moment, it's enough to know that this can be done within the Disks Manager program. You need to provide an empty folder via which the partition will be made available. Usually, under Linux, the `/mnt` directory is used for temporary mounts such as this, so click the Change button and browse to this directory. Alternatively, you can simply type `/mnt` into the Access Path field.
4. Click the Enable button, and then click the Browse button next to it to access your Windows file contents, as shown in Figure 12-7.

Note You should be able to read and write to FAT32 partitions, but you cannot write to or edit files in an NTFS partition. Although it is *technically* possible, it's not advisable because of various technical limitations. Therefore, the NTFS partition is made read-only. If you have a desperate need to *write* files to your Windows NTFS partition from within Ubuntu, you might consider converting it to FAT32. This can be done using a variety of commercial partitioning programs, such as PartitionMagic (www.symantec.com).



Figure 12-7. You can mount any partition, including those belonging to Windows, using the Disks Manager tool.

The Disks Manager program works for temporary access to your Windows files. If you want to make the mount of your NTFS partition permanent so it appears automatically every time you boot up, you must create a permanent mount point and also edit the `/etc/fstab` file. (The `fstab` file contains details of all your regular mounts and is discussed in Chapter 15.) However, before editing it, you'll need to create a folder that will act as the mount point for your Windows directory. Here are the steps:

1. Select Applications ► Accessories ► Terminal to open a GNOME Terminal window, and then type the following:

```
sudo mkdir /media/windows_ntfs
```

2. To open the `fstab` file in the Gedit text editor, type the following in the GNOME Terminal window:

```
sudo gedit /etc/fstab
```

3. You'll be asked to type your password, so do so.
4. In the Gedit window, scroll to the bottom of the file and add the following line (on a completely new line of its own):

```
/dev/hda1 /media/windows_ntfs ntfs nls=utf8,umask=0222 0 0
```

Type this exactly as it is written. You need to put two spaces between each element on the line. This step assumes that the Windows partition is the first on your hard disk. If you know it to be the second partition, change `/dev/hda1` to read `/dev/hda2`. But for most users who installed Ubuntu onto a hard disk that already contained Windows, the Windows partition will be `/dev/hda1`.

Following this, whenever you reboot your computer, the Windows partition should be available via an icon on the desktop and by clicking Places ► Computer.

Any files you copy from the NTFS partition will be read-only. To change this, right-click the file after you've copied it across, select Properties, click the Permissions tab, and put a check in the Write box on the Owner line.

Accessing Networked Files

The easiest way to access servers or shared folders on Windows workstations or servers over a network is to click Places ► Network Servers. This will start Nautilus and attempt to search for Windows machines on your local network, just like Network Neighborhood and My Network Places on the various versions of Windows. However, if you've ever used either of these services under Windows, you might also know how unreliable they can be—some computers simply don't appear in the list, others appear eventually after a wait, and others appear but then prove to be mysteriously inaccessible.

A far quicker and reliable method of accessing a Windows machine is to open a Nautilus file browser window (Places ► Home), and then click Go ► Location. In the box, type the following:

```
smb://Windows computer's IP address/
```

As you can see, you'll need to know the IP address of the Windows computer. Ask your system administrator or the person in charge of that machine for this address. You may also be prompted to enter a username and/or password to access the shared folder, as shown in Figure 12-8.

Note If you're accessing a Windows 95, 98, or Me shared folder, only password protection will have been set (these versions of Windows are unable to specify a username). However, when prompted by Nautilus, you still need to type *something* into the Username box to gain access—anything will do, as long as the password is correct. You cannot leave the Username box blank.

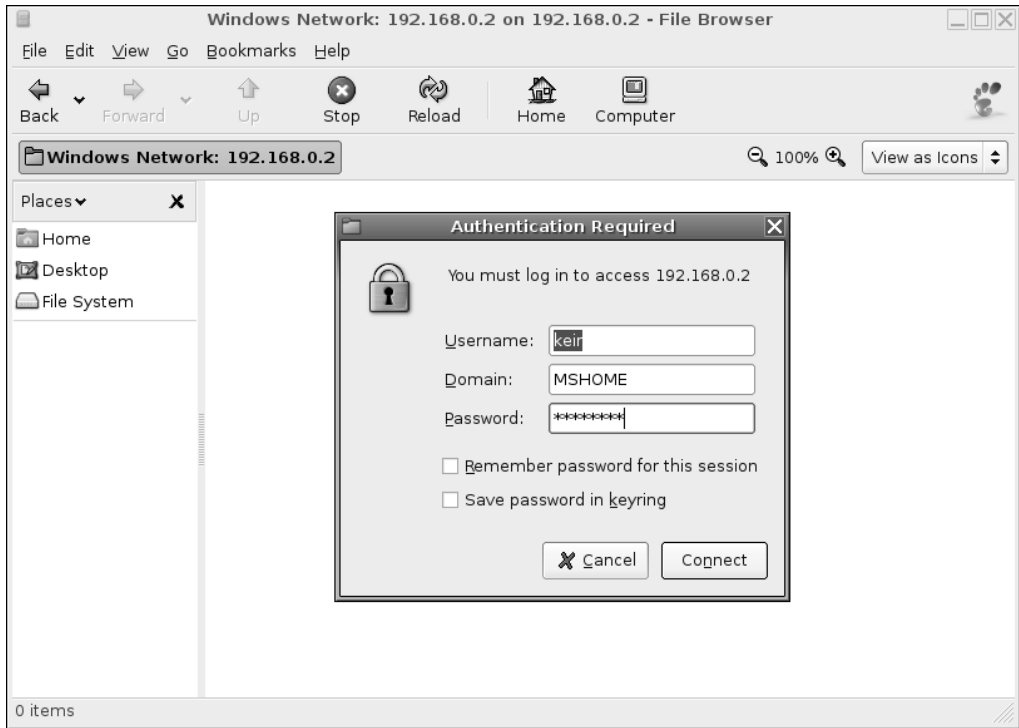


Figure 12-8. *If the share requires a username and/or password, you'll be invited to enter these.*

Accessing Removable Storage Devices

Ubuntu automatically makes available any CDs or DVDs you insert into your computer, and they'll appear instantly as icons on the desktop. The same is true of any card readers or USB memory devices that you use.

Alternatively, you can access the storage devices by clicking Places ► Computer. Here, you'll find icons for all of the storage devices attached to your computer, including the floppy disk drive, as shown in Figure 12-9. However, because of the way floppy disk drives work, Ubuntu isn't able to automatically detect if a floppy has been inserted. Instead, you'll need to double-click the icon, as with Windows.

Note In days of old, special tools were used to access MS-DOS floppies under Linux, and you might hear some Linux old hands talking about them. Nowadays, you can simply use Nautilus without needing to take any special steps.

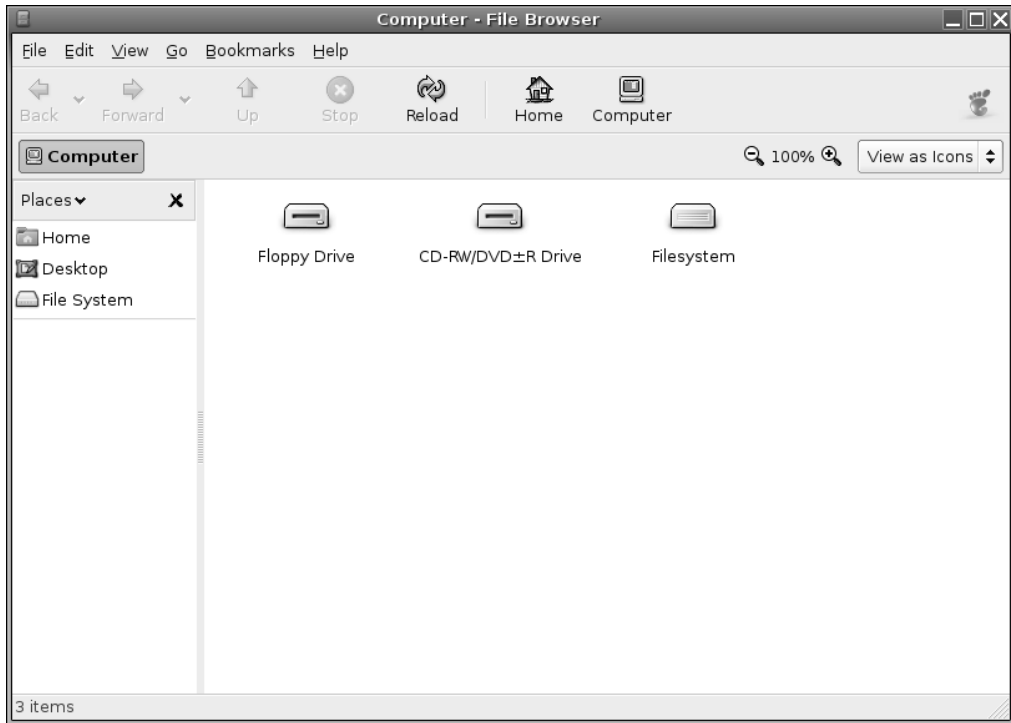


Figure 12-9. Select Places ► Computer to access your removable storage drives.

Whenever you double-click any entry in the Computer window, it will open a Nautilus file browser window. You can copy files by clicking and dragging, and right-clicking files offers virtually all the options you could need.

Tip You don't need to use Places ► Computer each time to access your floppy, CD, or DVD drive. These drives are mounted in the `/media` folder on your hard disk. Just browse to `/media/floppy`, and `/media/cdrom`.

Ejecting Media from Drives

Ubuntu isn't quite like Windows when it comes to ejecting or unplugging removable storage devices. In some cases, devices must be *unmounted*, which is to say that you need to tell Ubuntu that you're finished with the device in question and that you're about to unplug it.

In the case of CD or DVD discs, you can simply hit the Eject button on the drive itself. Ubuntu is clever enough to realize that the disc is being ejected, so it will automatically unmount the drive.

In the case of floppy disks, memory cards, and other USB storage devices, you'll need to right-click the icon and select Unmount Volume. Then you can then unplug and/or remove the device. This also applies when you're removing a memory card from a card reader—before pulling out the card from the reader, it needs to be unmounted.

Note It's necessary to close any files that were open on the device before unmounting, and even close any file browser windows that were accessing the device.

If you fail to unmount the device, Ubuntu will still believe the device is attached. This shouldn't cause too many problems, but it could crash any programs that were accessing the device.

Formatting Floppies

Formatting floppy disks isn't done in Nautilus. Instead, you must use a special program called Floppy Formatter. To start this program, click Applications ► System Tools ► Floppy Formatter. You'll see the program window shown in Figure 12-10.

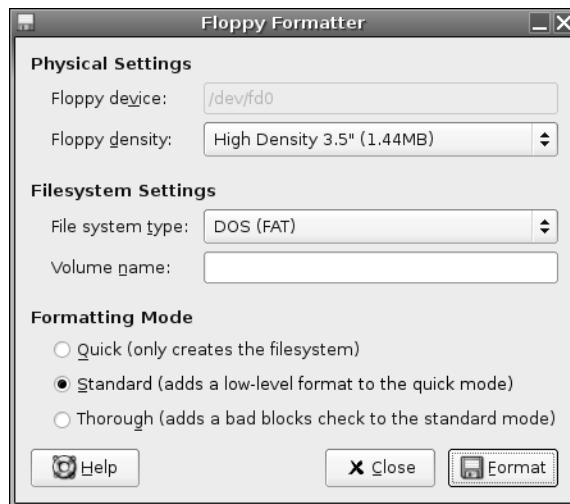


Figure 12-10. Formatting floppy disks is done using the Floppy Formatter tool.

Floppy Formatter is similar to the disk-formatting tool in Windows, and most of the options are self-explanatory. If you intend to share the disk with Windows users, make sure DOS (FAT) is selected in the File System Type box (it's possible to format a floppy using Ubuntu's own ext2 file system format, but there's little to be gained by doing so).

Summary

This chapter has led you on your first steps in exploring the Linux file system. The file system is vitally important to how Linux works, and we'll go into it in much depth in upcoming chapters.

Here, you were introduced to elementary concepts, such as where personal files are stored and the basic rules that govern what you can and cannot do with files. We also looked at the principle method of accessing files via the GUI: the Nautilus file manager. Additionally, you learned how to run programs manually, as well as how to access any Windows partition or files that may exist on your hard disk or across a network.

In Part 4 of this book, starting in the next chapter, we will look at some of the underlying technology that makes Ubuntu work and how you can gain more control over your computer. Chapter 13 introduces the BASH shell—perhaps the most powerful piece of software offered by Ubuntu to control your system.

PART 4



The Shell and Beyond



Introducing the BASH Shell

As you learned in Chapter 1, strictly speaking, the word *Linux* refers to just the kernel, which is the fundamental, invisible program that runs your PC and lets everything happen. However, on its own, the kernel is completely useless. It needs programs to let users interact with the PC and do cool stuff, and it needs a lot of system files (also referred to as *libraries*) to provide vital functions.

The GNU Project provides many of these low-level pieces of code and programs. This is why many people refer to the Linux operating system as GNU/Linux, giving credit to the fact that, without the GNU components, Linux wouldn't have gotten off the starting blocks.

The GNU Project provides various shell programs, too. Some of these offer graphical functionality, but most are text only. These text shell programs are also known as *terminal programs*, and they're often colloquially referred to as *command-line prompts*, in reference to the most important component they provide. This kind of shell lets you take control of your system in a quick and efficient way. Like a GUI, it's another way of interfacing with your computer, except that you type commands, rather than use a mouse.

By learning how to use the shell, you'll become the true master of your own system. In this part of the book, you'll learn all you need to know about using the shell. This chapter introduces the BASH shell, which is the default one in Ubuntu.

What Is the BASH Shell?

The best way of explaining the BASH shell to a Windows user is to compare it to the DOS command prompt. It lets you issue commands directly to the operating system via the keyboard without needing to mess around with the mouse and windows (although it is sometimes possible to use the mouse within a BASH shell to copy and paste text, and sometimes to control simple text-based menus). The big difference is that the BASH shell has commands for just about everything you might do on your system, whereas the DOS command prompt is restricted to tools capable of manipulating and viewing files and directories, and on Windows 2000/XP machines, configuring certain system settings.

In the old days, the DOS command prompt was also the visible layer of an entire operating system in which DOS programs were designed to be run. However, the shell is merely one of the many ways of accessing the Linux kernel and subsystems. It's true that many programs are designed to run via the BASH shell, but technically speaking, most actually run on the Linux operating system, and simply take input and show their input via the BASH shell.

Note Linux purists will point out another reason why the shell isn't exactly the same as a DOS command prompt within Windows: it doesn't run in virtual machine mode, a CPU trick by which part of the memory is subdivided to let programs run as if they had the PC all to themselves.

Linux finds itself with the BASH shell largely because Linux is a clone of Unix. In the early days of Unix, the text-based shell was all that was offered as a way of letting users control the computer. Typing commands in directly is one of the most fundamental ways of controlling any type of computer and, in the evolutionary scale, comes straight after needing to set switches and watch blinking lights in order to run programs.

That the BASH shell can trace its history back to the early days of Unix might sound like a tacit indication that the BASH is somehow primitive—far from it. It's one of the most efficient and immediate ways of working with your computer. Many people consider the command-line shell to be a way of using a computer that has yet to be superseded by a better method.

Note When you run a shell on a Linux system, the system refers to it as a `tty` device. This stands for teletypewriter, a direct reference to the old system of inputting data on what were effectively electronic typewriters connected to mainframe computers. These, in turn, took their names from the devices used to automate the sending and receiving of telegrams in the early part of the twentieth century.

Most Linux distributions come with a choice of different kinds of shell programs. However, the default shell is BASH, as is the case in Ubuntu. BASH stands for Bourne Again SHell. This is based on the Bourne shell, a tried-and-tested program that originated in the early days of Unix.

The other shells available include PDKSH (Public Domain Korn SHell, based on Korn Shell, another early Unix shell), and ZSH (Z SHell), a more recent addition. These are usually used by people who want to program Linux in various ways, or by those who simply aren't happy with BASH.

The BASH shell is considered by many to be the best of all worlds in that it's easy enough for beginners to learn, yet is able to grow with them and offer additional capabilities as necessary. BASH is capable of scripting, for example, which means you can even create your own simple programs.

Note Technically speaking, a *shell* refers to any type of user interface. The windowing system offered by Windows and Macintosh operating systems are a type of shell. However, many people in the Linux and Unix worlds use the word *shell* as shorthand for a shell that offers a command line.

Why Bother with the Shell?

You might have followed the instructions in Part 2 of this book and consider yourself an expert in Linux. But the real measure of a Linux user comes from your abilities at the shell.

In our modern age, the GUI is mistakenly considered “progress.” For instance, users of the Microsoft and Apple-based operating systems are quite used to using a mouse to navigate and perform various tasks. While it’s handy in certain situations—it would be difficult to imagine image editing without a mouse, for example—in many other situations, such as when manipulating files, directly typing commands is far more efficient.

Most modern Linux distributions prefer you to use the GUI to do nearly everything. This is because they acknowledge the dominance of Windows and realize they need to cater to mouse users who might not even know the shell exists (and, of course, programs like web browsers would be unusable without a GUI!). To this end, they provide GUI tools for just about every task you might wish to undertake. Ubuntu is strong in this regard, and you can configure a lot of things from the desktop.

However, it’s well worth developing at least some command-line shell skills, for a number of reasons:

It’s simple and fast. The shell is the simplest and fastest way of working with Ubuntu. As just one example, consider the task of changing the IP address of your network card. You could click the Systems menu, then the Administration option, then the Network option, and then double-click the entry in this list relating to your network card. That will take at least a minute or two if you know what you’re doing, and perhaps longer if it’s new to you. Alternatively, you could simply open a shell and type this:

```
ifconfig eth0 192.168.0.15 up
```

It’s versatile. Everything can be done via the shell—from deleting files, to configuring hardware, to creating MP3s. A lot of GUI programs actually make use of programs you can access via the shell.

It’s consistent among distributions. All Linux systems have shells and understand the same commands (broadly speaking). However, not all Linux systems will have Ubuntu’s graphical configuration programs. SUSE Linux uses its own GUI configuration tool, as does Mandrake Linux. Therefore, if you ever need to use another system, or decide to switch distributions, a reliance on GUI tools will mean learning everything from scratch. Knowing a few shell commands will help you get started instantly.

It’s crucial for troubleshooting. The shell offers a vital way of fixing your system should it go wrong. Your Linux installation might be damaged to the extent that it cannot boot to the GUI, but you’ll almost certainly be able to boot into a shell. A shell doesn’t require much of the system other than the ability to display characters on the screen and take input from the keyboard, which most PCs can do, even when they’re in a sorry state. This is why most rescue floppies offer shells to let you fix your system.

It’s useful for remote access. One handy thing about the shell is that you don’t need to be in front of your PC to use it. Programs like `ssh` let you log in to your PC across the Internet and use the shell to control your PC (as described in Chapter 34). For example, you can access data on a remote machine, or even fix it when you’re unable to attend the machine’s location. This is why Linux is preferred on many server systems when the system administrator isn’t always present on the site.

It's respected in the community. Using a shell earns you enormous brownie points when speaking to other Linux users. It separates the wheat from the chaff and the men from the boys (or women from the girls). If you intend to use Linux professionally, you will most certainly need to be a master at the shell.

Seen in this light, learning at least a handful of shell commands is vital to truly mastering your PC.

The drawback when using a command-line shell is that it's not entirely intuitive. Take for instance the command to change the network card's IP address:

```
ifconfig eth0 192.168.0.15 up
```

If you've never used the shell before, it might as well be Sanskrit. What on earth does `ifconfig` mean? And why is there the word `up` at the end?

Learning to use the shell involves learning terms like these. Hundreds of commands are available, but you really need to learn only around 10 or 20 for everyday use. The comparison with a new language is apt because, although you might think it daunting to learn new terminology, with a bit of practice, it will all become second nature. Once you've used a command a few times, you'll know how to use it in the future.

The main thing to realize is that the shell is your friend. It's there to help you get stuff done as quickly as possible. When you become familiar with it, you'll see that it is a beautiful concept. The shell is simple, elegant, and powerful.

When Should You Use the Shell?

The amount of use the Linux shell sees is highly dependent on the user. Some Linux buffs couldn't manage without it. They use it to read and compose e-mail, and even to browse the Web (usually using the Mutt and Lynx programs, respectively).

However, most people simply use it to manage files, view text files (like program documentation), and run programs. All kinds of programs—including GUI and command-line—can be started from the shell. As you'll learn in Chapter 29, unlike with Windows, installing a program on Ubuntu doesn't necessarily mean the program will automatically appear on the Applications menu. In fact, unless the installation routine is specifically made for the version of Linux you're running, this is unlikely. Therefore, using the shell is a necessity for most people.

Note Unlike with DOS programs, Ubuntu programs that describe themselves as “command-line” are rarely designed to run solely via the command-line shell. All programs are like machines that take input at one end and output objects at the other. Where the input comes from and where the output goes to is by no means limited to the command line. Usually, with a command-line program, the input and output are provided via the shell, and the programmer makes special dispensation for this, but this way of working is why GUI programs often make use of what might be considered shell programs. You'll often find that a GUI program designed to, for example, burn CDs, will also require the installation of a command-line program that will actually do the hard work for it.

There's another reason why the shell is used to run programs: you can specify how a particular program runs before starting it. For example, to launch the Totem Movie Player in full-screen mode playing the `myvideofile.mpg` file, you could type this:

```
totem --fullscreen myvideofile.mpg
```

This saves the bother of starting the program, loading a clip, and then selecting the full-screen option. After you've typed the command once or twice, you'll be able to remember it for the next time. No matter how much you love the mouse, you'll have to admit that this method of running programs is more efficient.

When you get used to using the shell, it's likely you'll have it open most of the time behind your other program windows.

Getting Started with the Shell

You can start the shell in a number of ways. The most common is to use a terminal emulator program. As its name suggests, this runs a shell inside a program window on your desktop.

You can start GNOME Terminal, the built-in GNOME shell emulator, by clicking Applications ► Accessories ► Terminal, as shown in Figure 13-1.

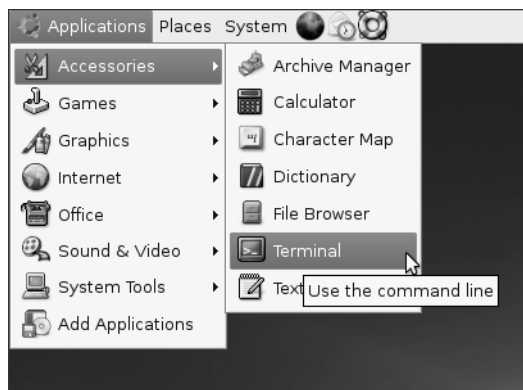


Figure 13-1. Start the The GNOME Terminal program from the Accessories submenu.

You'll see the terminal window—a blank, white window that's similar to a simple text editor window. It will show what's referred to as a *command prompt*: a few words followed by the dollar symbol: `$`. On my test system, this is what I see:

```
keir@ubuntu:~$
```

The first part is my username—the user account I created during installation and use to log in to the PC. After the `@` sign is the name of the PC, which I also chose when installing Ubuntu. The name of the PC isn't important on most desktop PCs; it's a legacy from the days of Unix. After the colon is the current directory you're browsing. In this example, the `~` symbol appears instead of an actual path or directory name. This is merely Linux shorthand for the user's home directory. In other words, wherever I see a `~` on my test PC, I read it as `/home/keir/`. After this is the dollar symbol (`$`), which indicates that I'm currently logged in as an ordinary

user, as opposed to the root user. However, unlike most other Linux distributions, Ubuntu doesn't use the root account during day-to-day operations, so this is a moot point. Finally, there is a cursor, and this is where you can start typing commands!

Note If you were to log in as root, a hash (#) would appear instead of the dollar prompt. This is important to remember because often in magazines and some computer manuals the use of the hash symbol before a command indicates that it should be run as root. In addition, if you use the rescue function of the install CD, you'll be running as root, and a hash will appear at the prompt. See Chapter 14 for more information about the root user.

Running Programs

When we refer to *commands* at the shell, we're actually talking about small programs. When you type a command to list a directory, for example, you're actually starting a small program that will do that job. Seen in this light, the shell's main function is to simply let you run programs—either those that are built into the shell, such as ones that let you manipulate files, or other, more complicated programs that you've installed yourself.

The shell is clever enough to know where your programs are likely to be stored. This information was given to it when you first installed Ubuntu and is stored in a system variable.

Note A *variable* is the method Linux uses to remember things like names, directory paths, or other data. There are many system variables that are vital for the running of Ubuntu.

The information about where your programs are stored, and therefore where Ubuntu should look for commands you type in, as well as any programs you might want to run, is stored in the PATH variable. You can take a look at what's currently stored there by typing the following:

```
echo $PATH
```

Don't forget that the difference between uppercase and lowercase letters matters to Ubuntu, unlike with Windows and DOS.

The echo command merely tells the shell to print something on screen. In this case, you're telling it to "echo" the PATH variable onto your screen. On my test PC, this returned the following information:

```
/usr/local/bin:/usr/local/sbin:/sbin:/usr/sbin:/bin:/usr/bin:/usr/bin/X11:/usr/games
```

Several directories are in this list, each separated by a colon.

Don't worry too much about the details right now. The important thing to know is that whenever you type a program name, the shell looks in each of the listed directories in sequence. In other words, when you type `ls`, the shell will look in each of the directories stored in the `PATH` variable, starting with the first in the list, to see if the `ls` program can be found. The first instance it finds is the one it will run. (The `ls` command gives you a directory listing, as described in the "Listing Files" section later in this chapter.)

But what if you want to run a program that is not contained in a directory listed in your `PATH`? In this case, you must tell the shell exactly where the program is. Here's an example:

```
/home/keir/myprogram
```

This will run a program called `myprogram` in the `/home/keir` directory. It will do this regardless of the directory you're currently browsing, and regardless of whether there is anything else on your system called `myprogram`.

If you're already in the directory where the program in question is located, you can type the following:

```
./myprogram
```

So, just enter a dot and a forward slash, followed by the program name.

Getting Help

Each command usually has help built in, which you can query (a little like typing `/?` after a command when using DOS). This will explain what the command does and how it should be used. In most cases, you'll see an example of the command in use, along with the range of command options that can be used with it. For example, you can get some instant help on the `ifconfig` command by typing this:

```
ifconfig --help
```

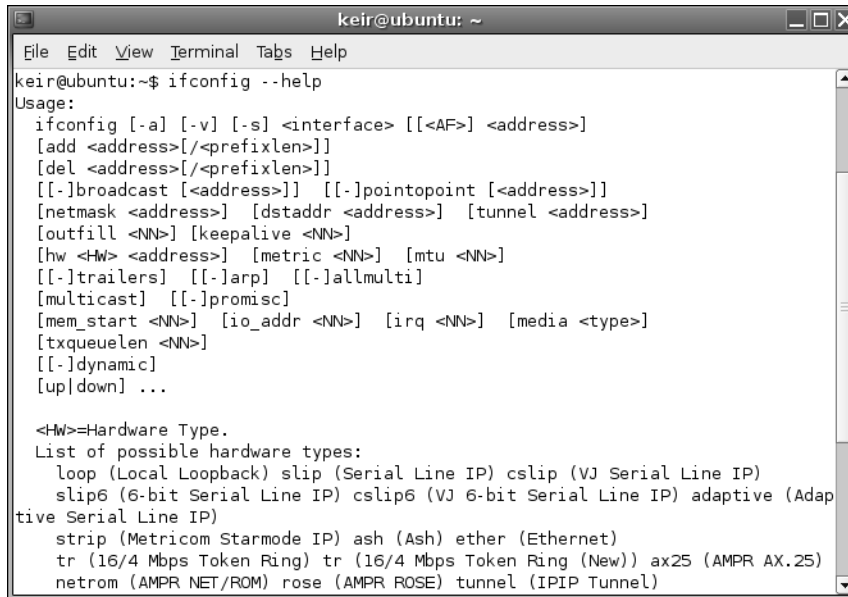
You'll see the help screen shown in Figure 13-2.

The `--help` option is fairly universal, and most programs will respond to it, although sometimes you might need to use a single dash. Just type the command along with `--help` to see what happens. You'll be told if you're doing anything wrong.

In addition, most commands have manuals that you can read to gain a fairly complete understanding of how they work. Virtually every Ubuntu setup has a set of these man pages, which can be accessed by typing this:

```
man <command>
```

However, man pages are often technical and designed for experienced Ubuntu users who understand the terminology.



```

keir@ubuntu:~$ ifconfig --help
Usage:
ifconfig [-a] [-v] [-s] <interface> [[<AF>] <address>]
[add <address>[/<prefixlen>]]
[del <address>[/<prefixlen>]]
[[-]broadcast [<address>]] [[-]pointopoint [<address>]]
[netmask <address>] [dstaddr <address>] [tunnel <address>]
[outfill <NN>] [keepalive <NN>]
[hw <HW> <address>] [metric <NN>] [mtu <NN>]
[[-]trailers] [[-]arp] [[-]allmulti]
[multicast] [[-]promisc]
[mem_start <NN>] [io_addr <NN>] [irq <NN>] [media <type>]
[txqueuelen <NN>]
[[-]dynamic]
[up|down] ...

<HW>=Hardware Type.
List of possible hardware types:
loop (Local Loopback) slip (Serial Line IP) cslip (VJ Serial Line IP)
slip6 (6-bit Serial Line IP) cslip6 (VJ 6-bit Serial Line IP) adaptive (Adap
tive Serial Line IP)
strip (Metricom Starmode IP) ash (Ash) ether (Ethernet)
tr (16/4 Mbps Token Ring) tr (16/4 Mbps Token Ring (New)) ax25 (AMPR AX.25)
netrom (AMPR NET/ROM) rose (AMPR ROSE) tunnel (IPIP Tunnel)

```

Figure 13-2. Most commands contain built-in help to give you a clue as to how they're used.

Some commands also have info pages, which offer slightly more down-to-earth guides. You can read these by typing this:

```
info <command>
```

If a command isn't covered by the info system, you'll be shown the default screen explaining basic facts about how the info command works.

Note that both man and info have their own man and info pages, explaining how they work. Just type man man or info info.

Running the Shell via a Virtual Console

As noted earlier, you can start the shell in a number of ways. The most common way among Linux diehards is via a virtual console. To access a virtual console, press Ctrl+Alt, and then press one of the function keys from F1 through F6 (the keys at the top of your keyboard).

Using a virtual console is a little like switching desks to a completely different PC. Pressing Ctrl+Alt+F1 will cause your GUI to disappear, and the screen to be taken over by a command-line prompt (don't worry; your GUI is still there and running in the background). You'll be asked to enter your username and your password.

Any programs you run in a virtual console won't affect the rest of the system, unless they're system commands. (As discussed in Chapter 16, one way to rescue a crashed GUI program is to switch to a virtual console and attempt to terminate the program from there.)

You can switch back to the GUI by pressing Ctrl+Alt+F7. Don't forget to quit your virtual console when you're finished with it, by typing exit.

BOOTING INTO THE SHELL

If you're really in love with the shell, you can choose to boot into it, avoiding the GUI completely (although you can later start the GUI by typing `startx` at the command line).

Booting into the shell is done by defining a custom run level. A *run level* is how the operating mode that Ubuntu is currently running in is described. For example, one particular run level might start a GUI, while another might start only a command prompt.

There are usually seven run levels under Linux, numbered from 0 to 6. Not all of them do something interesting. On Ubuntu, run levels 2 through 5 are all the same. Each runs the GUI. Run level 1 runs a command prompt, so it might seem ideal for booting into the shell, but it also shuts down a few essential services. This means it isn't suitable for day-to-day use.

The trick is to take one of the existing run levels and alter it slightly so that it doesn't run a GUI by default. On many distributions, run level 3 is reserved for this purpose, so it makes sense to alter it under Ubuntu. (For what it's worth, the default Ubuntu run level is 2.)

Stopping Ubuntu from running a GUI upon booting is simply a matter of stopping the program that appears when Ubuntu boots—GDM. This provides the login window that appears and starts the whole graphical subsystem. Type the following command at the shell to remove the shortcut to GDM within the run level 3 configuration:

```
sudo rm /etc/rc3.d/S13GDM
```

After this, you'll need to tell Ubuntu to boot straight to run level 3, rather than the default of 2. You do this by editing the `/etc/inittab` file, which is one of the first configuration files Ubuntu reads when booting. Issue the following command at the shell to open the file in the Gedit text editor:

```
sudo gedit /etc/inittab
```

Then look for the following line, which will be near the top of the file:

```
id:2:initdefault
```

And edit it so that it reads:

```
id:3:initdefault
```

Then save the file. From now on, you'll always boot straight to a BASH prompt. To restore things to the way they were, simply restore the line in the `/etc/inittab` file to the way it appeared originally.

Working with Files

So let's start actually using the shell. If you've ever used DOS, then you have a head start over most shell beginners, although you'll still need to learn some new commands. Table 13-1 shows various DOS commands alongside their Ubuntu equivalents. This table also serves as a handy guide to some BASH commands, even if you've never used DOS. In Appendix B, you'll

find a comprehensive list of useful shell commands, together with explanations of what they do and examples of typical usage.

Table 13-1. *DOS Commands and Their Shell Equivalents*

Command	DOS Command	Linux Shell Command	Usage
Copy files	COPY	cp	cp <filename> <new location>
Move files	MOVE	mv	mv <filename> <new location>
Rename files	RENAME	mv	mv <old filename> <new filename> ¹
Delete files	DEL	rm	rm <filename> ²
Create directories	MKDIR	mkdir	mkdir <directory name>
Delete directories	DELTREE/RMDIR	rm	rm -rf <directory name>
Change directory	CD	cd	cd <directory name>
Edit text files	EDIT	vi	vi <filename>
View text files	TYPE	less	less <filename> ³
Print text files	PRINT	lpr	lpr <filename>
Compare files	FC	diff	diff <file1> <file2>
Find files	FIND	find	find -name <filename>
Check disk integrity	SCANDISK	fsck	fsck ⁴
View network settings	IPCONFIG	ifconfig	ifconfig
Check a network connection	PING	ping	ping <address>
View a network route	TRACERT	tracert	tracert <address>
Clear screen	CLS	clear	clear
Get help	HELP	man	man <command> ⁵
Quit	EXIT	exit	exit

¹ The BASH shell offers a rename command, but this is chiefly used to rename many files at once.

² To avoid being asked to confirm each file deletion, you can add the -f option. Be aware that the rm command deletes data instantly, without the safety net of the Recycle Bin, as with the GNOME desktop.

³ Use the cursor keys to move up and down in the document. Type Q to quit.

⁴ This is a system command and can be run only on a disk that isn't currently in use. To scan the main partition, you'll need to boot from the installation CD and select the rescue option. Then issue the fsck command.

⁵ The info command can also be used.

CREATING ALIASES

If you've ever used DOS, you might find yourself inadvertently typing DOS commands at the shell prompt. Some of these will actually work, because most distribution companies create command aliases to ease the transition of newcomers to Linux.

Aliases mean that whenever you type certain words, they will be interpreted as meaning something else. However, an alias won't work with any of the command-line switches used in DOS. In the long run, you should try to learn the BASH equivalents.

You can create your own command aliases quickly and simply. Just start a BASH shell and type the following:

```
alias <DOS command>='<Linux shell command>'
```

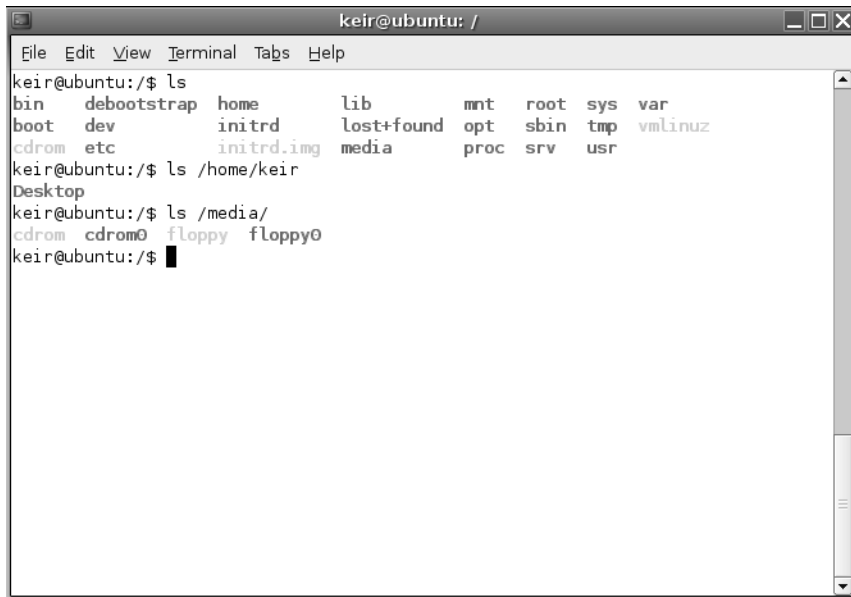
For example, to create an alias that lets you type `del` instead of `rm`, type this:

```
alias del='rm'
```

Note that the Ubuntu command must appear in single quotation marks.

Listing Files

Possibly the most fundamentally useful BASH command is `ls`. This will list the files in the current directory, as shown in Figure 13-3. If you have a lot of files, they might scroll off the screen. If you're running GNOME Terminal, you can use the scroll bar on the right side of the window to view the list.

A screenshot of a terminal window titled "keir@ubuntu: /". The window has a menu bar with "File", "Edit", "View", "Terminal", "Tabs", and "Help". The terminal shows the following text:

```
keir@ubuntu:/$ ls
bin  debootstrap  home      lib      mnt  root  sys  var
boot dev          initrd   lost+found  opt  sbin  tmp  vmlinuz
cdrom etc          initrd.img  media    proc  srv  usr

keir@ubuntu:/$ ls /home/keir
Desktop

keir@ubuntu:/$ ls /media/
cdrom  cdrom0  floppy  floppy0

keir@ubuntu:/$ █
```

Figure 13-3. The `ls` command lists the files in the current directory.

Having the files scroll off the screen can be annoying, so you can cram as many as possible onto each line by typing the following:

```
ls -m
```

The dash after the command indicates that you're using a command option. These are also referred to as command-line *flags* or *switches*. Nearly all shell commands have options like this. In fact, some commands won't do anything unless you specify various options. In the case of the `ls` command, only one dash is necessary, but some commands need two dashes to indicate an option.

You can see a list of all the command options for `ls` by typing the following (ironically, itself a command option):

```
ls --help
```

Once again, the output will scroll off the screen, and you can use the window's scroll bars to examine it. (In Chapter 17, you'll learn a trick you can use to be able to read this output without needing to fiddle around with the scroll bars, even if there's screen after screen of it.)

With most commands, you can use many command options at once, as long as they don't contradict each other. For example, you could type the following:

```
ls -lh
```

This tells the `ls` command to produce "long" output and also to produce "human-readable" output. The long option (`-l`) lists file sizes and ownership permissions, among other details (permissions are covered in the next chapter). The human-readable option (`-h`) means that rather than listing files in terms of bytes (such as 1029725 bytes), it will list them in kilobytes and megabytes. Notice that you can simply list the options after the dash; you don't need to give each option its own dash.

Caution I've said it before, and I'll say it again: don't forget that case-sensitivity is vitally important in Ubuntu! Typing `ls -L` is not the same as typing `ls -l`. It will produce different results.

Copying Files

Another useful command for dealing with files is `cp`, which copies files. You can use the `cp` command in the following way:

```
cp myfile /home/keir/
```

This will copy the file to the location specified.

One important command-line option for `cp` is `-r`. This stands for recursive and tells BASH that you want to copy a directory and its contents (as well as any directories within this directory). Most commands that deal with files have a recursive option.

Note Only a handful of BASH commands default to recursive copying. Even though it's extremely common to copy folders, you still need to specify the `-r` command option most of the time.

One curious trick is that you can copy a file from one place to another but, by specifying a filename in the destination part of the command, change its name. Here's an example:

```
cp myfile /home/keir/myfile2
```

This will copy `myfile` to `/home/keir`, but rename it as `myfile2`. Be careful not to add a final slash to the command when you do this. In the example here, doing so would cause BASH to think that `myfile2` is a directory.

This way of copying files is a handy way of duplicating files. By not specifying a new location in the destination part of the command, but still specifying a different filename, you effectively duplicate the file within the same directory:

```
cp myfile myfile2
```

This will result in two identical files: one called `myfile` and one called `myfile2`.

Moving Files

The `mv` command is similar to `cp`, except that rather than copying the file, the old one is removed. You can move files from one directory to another, for example, like this:

```
mv myfile /home/keir/
```

You can also use the `mv` command to quickly rename files:

```
mv myfile myfile2
```

Figure 13-4 shows the results of using `mv` to rename a file.

Note Getting technical for a moment, moving a file in Linux isn't the same as in Windows, where a file is copied and then the original deleted. Under Ubuntu, the file's absolute path is rewritten, causing it to simply appear in a different place in the file structure. However, the end result is the same.

A terminal window titled "keir@ubuntu: ~/mydirectory" with a menu bar (File, Edit, View, Terminal, Tabs, Help). The terminal shows the following sequence of commands and output:

```
keir@ubuntu:~/mydirectory$ ls
myfile
keir@ubuntu:~/mydirectory$ mv myfile myfile2
keir@ubuntu:~/mydirectory$ ls
myfile2
keir@ubuntu:~/mydirectory$
```

Figure 13-4. You can also use the `mv` command to rename files.

Deleting Files

But how do you get rid of files? Again, this is relatively easy, but first a word of caution: the shell doesn't operate any kind of Recycle Bin. Once a file is deleted, it's gone forever. (There are utilities you can use to recover files, but these are specialized tools and aren't to be relied on for day-to-day use.)

Removing a file is achieved by typing something like this:

```
rm myfile
```

It's as simple as that.

You'll be asked to confirm the deletion after you issue the command. If you want to delete a file without being asked to confirm it, type the following:

```
rm -f myfile
```

The `f` stands for force (that is, force the deletion).

If you try to use the `rm` command to remove a directory, you'll see an error message. This is because the command needs an additional option:

```
rm -rf mydirectory
```

As noted earlier, the `r` stands for recursive and indicates that any folder specified afterwards should be deleted, in addition to any files it contains.

Tip You might have used wildcards within Windows and DOS. They can be used within Ubuntu, too. For example, the asterisk (*) can be used to mean any file. So, you can type `rm -f *` to delete all files within a directory, or type `rm -f myfile*` to delete all files that start with the word `myfile`. But remember to be careful with the `rm` command. Keep in mind that you cannot salvage files easily if you accidentally delete them!

Changing and Creating Directories

Another handy command is `cd`, for change directory. This lets you move around the file system, from directory to directory. Say you're in a directory that has another directory in it, named `mydirectory2`. Switching to it is easy:

```
cd mydirectory2
```

But how do you get out of this directory once you're in it? Try the following command:

```
cd ..
```

The `..` refers to the “parent” directory, which is the one containing the directory you're currently browsing. Using two dots to indicate this may seem odd, but it's just the way that Ubuntu (and Unix before it) does things. It's one of the many conventions that Unix relies on and that you'll pick up as you go along.

You can create directories with the `mkdir` command:

```
mkdir mydirectory
```

Summary

This chapter introduced the command-line shell, considered by many to be the heart of Linux. We've discussed its similarities to MS-DOS, and shown that these are only cursory; knowledge of DOS doesn't equate to skill within BASH. In the long run, you should work to polish your BASH skills.

This chapter also introduced some elementary commands used within BASH, such as those used to provide directory listings and to copy files. We looked at how you can use command-line options to control BASH tools. In many cases, these are mandatory, so you learned how the BASH shell itself can be used to investigate a command and find out vital information about how it works.

At this point, your newfound knowledge will have no doubt caused you to venture into the Ubuntu file system itself, which can be a confusing, if not terrifying, place for the inexperienced. But don't worry. The next chapter explains everything you need to know about the file system and what you'll find in it.



Understanding Linux Files and Users

Most of us are used to dealing with files—the things that live on our hard disks, floppies, and CD-ROMs, and contain data and program code. It should come as no surprise that Linux has its own file structure, which is different from Windows, in terms of where data is stored and also the underlying technology.

Taking a page from Unix, Ubuntu takes the concept of the file system to an extreme. To Ubuntu, almost everything is treated as a file: your PC's hardware, network computers connected to your PC, information about the current state of your computer . . . almost everything finds a home within the Linux file system.

Linux places an equal emphasis on the users of the system. They own the various files and can decide who can and cannot access various files they create.

In this chapter, we'll delve into users, files, and permissions. You'll be introduced to how Ubuntu handles files and how files are tied into the system of user accounts.

Real Files and Virtual Files

Linux sees virtually everything as a series of files. This might sound absurd and certainly requires further explanation.

Let's start with the example of plugging in a piece of hardware. Whenever you attach something to a USB socket, the Linux kernel finds it, sees if it can make the hardware work, and, if everything checks out okay, it will usually make the hardware available as a file under the `/dev` directory on your hard disk (`dev` is short for devices). Figure 14-1 shows an example of a `/dev` directory.

The file created in the `/dev` directory is not a real file, of course. It's a file system shortcut plumbed through to the input and output components of the hardware you've just attached.

Note As a user, you're not expected to delve into the `/dev` directory and deal with this hardware directly. Most of the time, you'll use various software packages that will access the hardware for you, or use special BASH commands or GUI programs to make the hardware available in a more accessible way for day-to-day use.

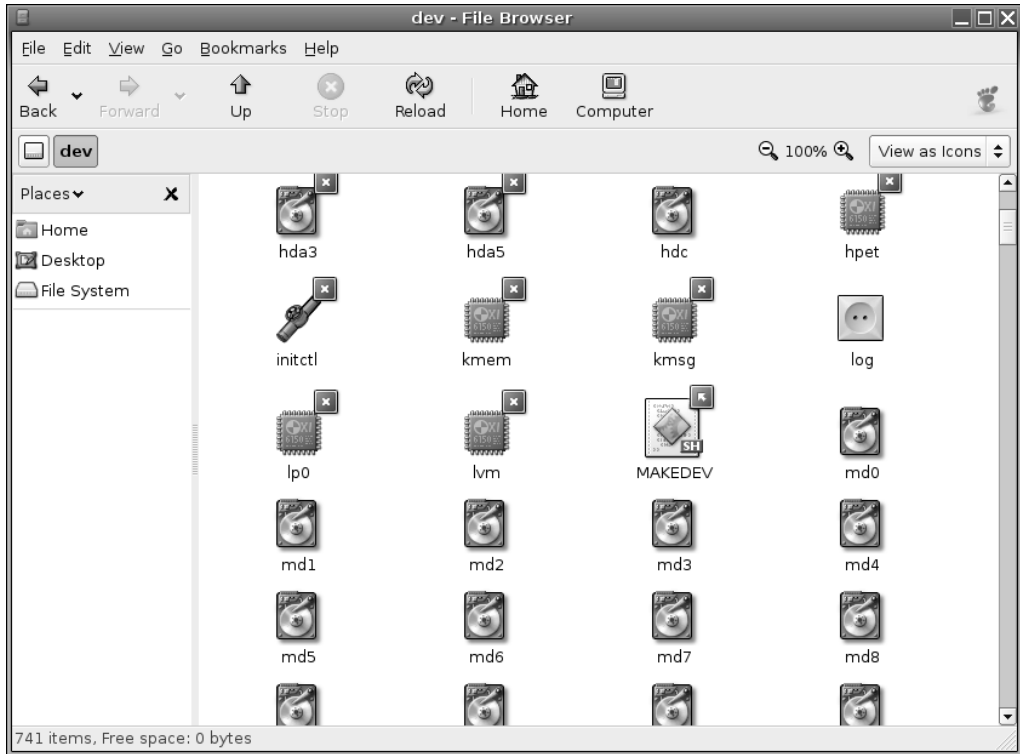


Figure 14-1. Hardware devices under Linux are accessed as if they were files and can be found in the `/dev` folder.

Here’s another example. Say you’re working in an office and you want to connect to a central file server. To do this under Linux, you must *mount* the files that the server offers, making it a part of the Ubuntu file system. Doing this involves creating an empty directory (or using one that already exists) and using the `mount` command at the BASH shell to make the server’s contents magically appear whenever that directory is accessed. We’ll discuss how this is done later in this chapter, in the “Mounting” section (but remember that Ubuntu may do this automatically, as discussed in Chapter 12).

Once the network server is mounted, it is treated exactly like a directory on your hard disk. You can copy files to and from it, just as you would normally, using the same tools as you use for dealing with any other files. In fact, less knowledgeable users won’t even be aware that they’re accessing something that isn’t located on their PC’s hard disk (or, technically speaking, within their Ubuntu partition and file system).

By treating everything as a file, Linux makes system administration easier. To probe and test your hardware, for example, you can use the same tools you use to manipulate files.

So how do you know which files are real and which are virtual? One method is to use the following command, which was introduced in the previous chapter:

```
ls -l
```

The `-l` option tells the `ls` command to list nearly all the details about the files. If you do this in GNOME Terminal, you'll see that the listing is color-coded. Table 14-1 shows what each color indicates. The command returns a lot of additional information, including who owns which file and what you and others can do with it. This requires an understanding of users and file permissions, which we'll discuss next.

Tip The command `ls -la` will give you even more information—perhaps too much for general use. In most instances, `ls -l` should show enough information.

Table 14-1. *Color-Coding Within GNOME Terminal*

Color	Type of File
Black text	Standard file
Light-blue text	Directory
Black outline with yellow text	Virtual device ¹
Green text	Program or script ²
Cyan text	Symbolic link to another file ³
Pink text	Image file
Red text	Archive ⁴

¹ This is found only in the `/dev` directory.

² Technically speaking, green text indicates a program or script that has merely been marked as being executable.

³ This is similar to a Windows desktop shortcut.

⁴ Installation files are also marked red because they're usually contained in archives.

Users and File Permissions

The concept of users and permissions is as important to Ubuntu as the idea of a central and all-encompassing file system. In fact, the two are implicitly linked.

When initially installing Linux, you should have created at least one user account. By now, this will have formed the day-to-day login that you use to access Linux and run programs.

Although you might not realize it, as a user, you also belong to a group. In fact, every user on the system belongs to a group. Under Ubuntu, ordinary users belong to a group based on their username (under other versions of Linux, you might find that you belong to a group called `users`).

Note Groups are yet another reminder of Ubuntu’s Unix origins. Unix is often used on huge computer systems with hundreds or thousands of users. By putting each user into a group, the system administrator’s job is a lot easier. When controlling system resources, the administrator can control groups of users rather than hundreds of individual users. On most home user PCs, the concept of groups is a little redundant, because there’s normally a single user, or at most, two or three. However, the concept of groups is central to the way that Linux handles files.

A standard user account under Ubuntu is normally limited in what it can do. As a standard user, you can save files to your own private area of the disk, located in the /home directory, as shown in Figure 14-2, but usually nowhere else. You can move around the file system, but some directories are strictly out of bounds. In a similar way, some files can be opened as read-only, so you cannot save changes to them. All of this is achieved using file permissions.

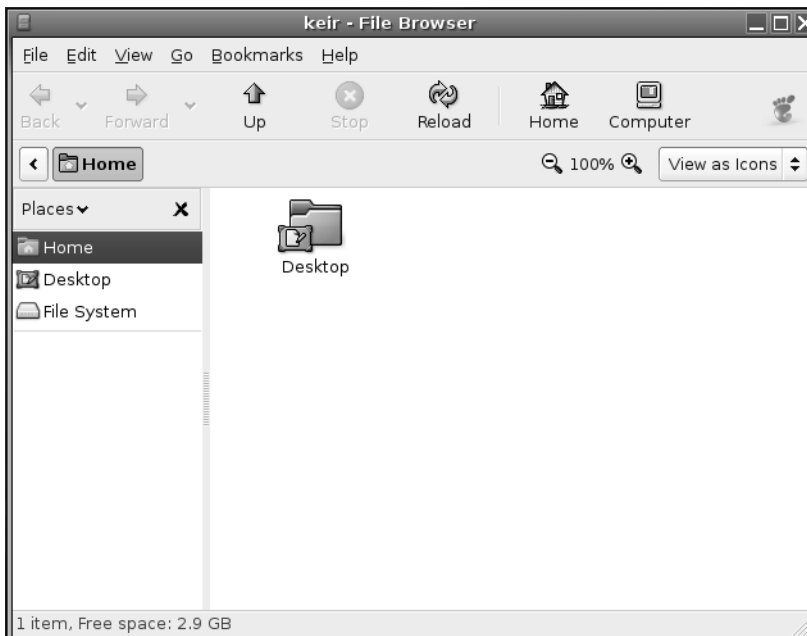


Figure 14-2. Your personal directory within home is your area on the hard disk. This is enforced via file permissions.

Every file and directory is owned by a user. In addition, files and directories have three separate settings that indicate who within the Linux system can read them, who can write to them, and, if the file in question is “runnable” (usually a program or a script), who can run it (“execute” it). In the case of directories, it’s also possible to set who can browse them, as well as who can write files to them. If you try to access a file or directory for which you don’t have permission, you’ll be turned away with an “access denied” error message.

ROOT VS. SUDO

Most versions of Linux have two types of user accounts: standard and root. Standard users are those who can run programs on the system but are limited in what they can do. Root users have complete run of the system, and as such, are often referred to as “superusers.” They can access and/or delete whatever files they want. They can configure hardware, change settings, and so on.

Most versions of Linux create a user account called root and let users log in as root to perform system maintenance. However, for practical as well as security reasons, most of the time the user is logged in as a standard user.

Ubuntu is different in that it does away with the root account. Instead, it allows certain users, including the one created during installation, to temporarily adopt root-like powers. You will already have encountered this when configuring hardware. As you’ve seen, all you need to do is type your password when prompted in order to administer the system.

This way of working is referred to as *sudo*, which is short for “superuser do.” In fact, the command `sudo` will let you adopt root powers at the shell prompt—simply preface any command with `sudo` in order to run it with root privileges.

In some ways, the `sudo` system is slightly less secure than using a standard root account. But it’s also a lot simpler. It reduces the chance of serious errors, too. Any command or tweak that can cause damage will invariably require administrative powers, and therefore require you to type your password or preface the command with `sudo`. This serves as a warning and prevents mistakes.

If you’re an experienced Linux user and want to invoke the root account, simply type the following at the command prompt:

```
sudo passwd root
```

Then type a password. If you subsequently want to deactivate the root account, type this:

```
sudo passwd -l root
```

Viewing Permissions

When you issue the `ls -l` command, each file is listed on an individual line. Here’s an example of one line of a file listing from my test PC:

```
-rw-r--r--  2 keir keir 673985982 2005-10-31 17:19 myfile
```

The `r`, `w`, and `-` symbols on the very left of the listing indicate the file permissions. The permission list usually consists of the characters `r` (for read), `w` (for write), `x` (for execute), or `-` (meaning none are applicable).

They’re followed by a number indicating the link count, which you can ignore. After this is listed the owner of the file (`keir` in the example) and the group that he belongs to (`keir`). This is followed by the file size (in bytes), then the date and time the file was last accessed, and finally the filename itself appears.

The file permissions part of the listing might look confusing, but it's actually quite simple. To understand what's going on, you need to split it into groups of four, as illustrated in Figure 14-3.




Figure 14-3. The file permissions part of a file listing can be broken down into four separate parts.

The four groups are as follows:

- **Type of file:** This character represents the file type. A standard data file is indicated with a dash (-). Most files on your system fall into this category. A `d` shows that the entry is not a file, but a directory. Table 14-2 lists the file type codes.
- **User permissions:** Next come the permissions of the person who owns the file. The three characters indicate what the person who owns the file can do with it. The owner of a file is usually the user who created it, although it's also possible to change the owner later on. In this example, you see `rw-`. This means that the owner of the file can read (`r`) and write (`w`) the file. In other words, he can look at it and also save changes to it. However, there's a dash afterwards, and this indicates that the user cannot execute the file. If this were possible, there would be an `x` in this spot instead.
- **Group permissions:** After the owner's permissions are the permissions given to members of that user's group. This is indicated by another three characters in the same style as those for user permissions. In the example, the group's permission is `r--`, which means that the group members can read the file but don't have permission to write to it, since there's a dash where the `w` would normally appear. In other words, as far as they're concerned, the file is read-only.
- **Everyone else's permissions:** The last set of permissions indicates the permissions of everyone else on the system (other users in other groups). In the example, they can only read the file (`r`); the two dashes afterwards indicate that they cannot write to the file nor execute it.

Table 14-2. File Type Codes

Code	File Type
-	Standard file
d	Standard directory
l	Symbolic link (a shortcut to another file)

Table 14-2. *File Type Codes*

Code	File Type
p	Named pipe (a file that acts as a conduit for data between two programs)
s	Socket (a file designed to send and receive data over a network)
c	Character device (a hardware device driver, usually found in /dev)
b	Block device (a hardware device driver, usually found in /dev)

As you might remember from Windows, programs are stored as files on your hard disk, just like standard data files. On Linux, program files need to be explicitly marked as being executable. This is indicated in the permission listing by an `x`. Therefore, if there's no `x` in a file's permissions, it's a good bet that the file in question isn't a program or script (although this isn't always true for various technical reasons).

To make matters a little more confusing, if the entry in the list of files is a directory (indicated by a `d`), then the rules are different. In this case, an `x` indicates that the user can access that directory. If there's no `x`, then the user's attempts to browse to that directory will be met with an "access denied" message.

File permissions can be difficult to understand, so let's look at a few real-world examples. These examples assume that you're logged in to Linux as the user `keir`.

LESS COMMON FILE TYPES

Instead of the `x` in the list of permissions for a directory, you might sometimes see a `t`. This means that the only people who can delete or alter a file in that directory are the users who created the file in the first place. This is a useful option to have in some circumstances.

You might sometimes see a set of permissions like `rws`. The `s` stands for `setuid`. Like `x`, it indicates that the file is executable, except, in this case, it means that the file will be run with the permissions of the person who owns it, rather than the user who is executing it. In other words, if user `frank` tries to run a program owned by `keir` that has the execute permission set as `s`, that program will be run as if `keir` were running it. This is very useful, because it can make programs that require root powers usable by ordinary users, although this brings with it obvious security risks.

Typical Data File Permissions

Here's the first example:

```
-rw-rw---- 2 keir keir 1450 2004-07-07 09:19 myfile2
```

You see immediately that this file is owned by user `keir` because that username appears directly after the permissions. You also see that this user is a member of the group `keir`.

Reading the file permissions from left to right, you see that the initial character is a dash. That indicates that this is an ordinary file and has no special characteristics. It's also not a directory.

After that is the first part of the permissions, `rw-`. These are the permissions for owner of the file, `keir`. You're logged in as that user, so this file belongs to you, and these permissions apply to you. You can read and write the file, but not execute it. Because you cannot execute the file, you can infer that this is a data file rather than a program (there are certain exceptions to this rule, but we'll ignore them for the sake of simplicity).

Following this is the next part of the file permissions, `rw-`. This tells you what other members of your group can do with the file. It's fairly useless information if you're the only user of your PC but, for the record, you're told that anyone else belonging to the group called `keir` can also read and write the file, but not execute it. If you're not the only user of a computer, group permissions can be important. The "Altering Permissions" section, coming up shortly, describes how to change file permissions to control who can access files.

Finally, the last three characters tell you the permissions of everyone else on the system. The three dashes (`---`) mean that they have no permissions at all regarding the file. There's a dash where the `r` normally appears, so they cannot even read it. The dashes afterwards tell you they cannot write to the file or execute it. If they try to do anything with the file, they'll get a "permission denied" error.

Permissions on a User's Directory

Here's example number two:

```
drwxr-xr-x  7 keir keir 824 2004-07-07 10:01 mydirectory
```

The list of permissions starts with `d`, which tells you that this isn't a file but a directory. After this is the list of permissions for the owner of the directory (`keir`), who can read files in the directory and also create new ones there. The `x` indicates that you can access this directory, as opposed to being turned away with an "access denied" message. You might think being able to access the directory is taken for granted if the user can read and write to it, but that's not the case.

Next are the permissions for the group members. They can read files in the directory but not write any new ones there (although they can modify files already in there, provided the permissions of the individual files allow this). Once again, there's an `x` at the end of their particular permission listing, which indicates that the group members can access the directory.

Following the group's permissions are those of everyone else. They can read the directory and browse it, but not write new files to it, as with the group users' permissions.

Permissions on a Directory Owned by Root

Here's the last example:

```
drwx----- 25 root root 1000 2004-08-06 15:44 root
```

You can see that the file is owned by root. Remember that in this example, you're logged in as keir and your group is keir.

The list of permissions starts with a d, so you can tell that this is actually a directory. After this, you see that the owner of the directory, root, has permission to read, write, and access the directory.

Next are the permissions for the group: three dashes. In other words, members of the group called root have no permission to access this directory in any way. They cannot browse it, create new files in it, or even access it.

Following this are the permissions for the rest of the users. This includes you, because you're not the user root and don't belong to its group. The three dashes means you don't have permission to read, write, or access this directory.

SWITCHING USERS

If you have more than one user set up on your system, it's possible to switch users on the fly while you're working at the shell. On my test PC, I have an additional user account called frank. While logged in as any user, I can temporarily switch to this user by typing the following command, which stands for substitute user:

```
su frank
```

I'll then be asked for user frank's password. Once this is typed, I will effectively have logged in as user frank. Any files I create will be saved with frank's ownership.

If you created a root account (by using the command `sudo passwd root`), you can temporarily switch into it by typing just `su`, without any username afterwards.

To return to your own account from any other account, type `exit`.

Altering Permissions

You can easily change permissions of files and directories by using the `chmod` command. For example, if you want to change a file so that everyone on the system can read and write to it, type the following:

```
chmod a+rw myfile
```

In other words, you're adding add read and write (`rw`) permissions for all users (`a`). Here's another example:

```
chmod a-w myfile
```

This tells Linux that you want to take away (-) the ability of all users (`a`) to write (`w`) to the file. However, you want to leave the other permissions as they are. You can substitute a with `g` to change group permissions instead.

The most useful use of `chmod` is in making a program file that you've downloaded executable. Due to the way the Internet works, if you download a program to install on your computer, it can lose its executable status while in transit. In this case, issue the following command:

```
chmod +x myprogram
```

Because nothing is specified before the `+x`, the shell assumes that the changes to be applied to the file are for the current user only.

To change the owner of a file, use the `chown` command. For security reasons this must be prefaced with the `sudo` command, which is to say that `chown` (and `chgrp`, to change the group), require superuser powers.

For example, to set the owner of `myfile` as `frank`, type this command:

```
sudo chown frank myfile
```

You can also change the owner *and* the group of a file using `chown`. Simply type each separated by a period:

```
sudo chown frank.mygroup myfile
```

This will change `myfile` so that its owner is `frank` and its group is `mygroup`.

To change the group of a file, you can use the `chgrp` command in exactly the same way as `chown`:

```
sudo chgrp mygroup myfile
```

The File System Explained

Now that you understand the principles of files and users, we can take a bird's-eye view of the Linux file system and start to make sense of it.

You might already have ventured beyond the `/home` directory and wandered through the file system. You no doubt found it thoroughly confusing, largely because it's not like anything you're used to. The good news is that it's not actually very hard to understand. If nothing else, you should be aware that nearly everything can be ignored during everyday use.

Note The Ubuntu file system is referred to as a *hierarchical* file system. This means that it consists of a lot of directories that contain files. Windows also uses a hierarchical file system. Ubuntu refers to the very bottom level of the file system as the root. This has no connection with the root user.

You can switch to the root of the file system by typing the following shell command:

```
cd /
```

When used on its own, the forward slash is interpreted as a shortcut for root.

If I do this on my PC, and then ask for a long file listing (`ls -l`), I see the following:

```

total 112
drwxr-xr-x  2  root  root  4096 2005-10-14 22:29  bin
drwxr-xr-x  3  root  root  4096 2005-10-14 18:25  boot
lrwxrwxrwx  1  root  root   11 2005-10-14 18:01  cdrom -> media/cdrom
drwxr-xr-x  2  root  root  4096 2005-10-14 18:09  debootstrap
drwxr-xr-x 12  root  root   0 2005-10-16 12:25  dev
drwxr-xr-x 100 root  root  4096 2005-10-16 12:25  etc
drwxr-xr-x  3  root  root  4096 2005-10-14 18:20  home
drwxr-xr-x  2  root  root  4096 2005-10-14 18:02  initrd
lrwxrwxrwx  1  root  root   28 2005-10-14 18:08  initrd.img -> boot/initrd.im
g-2.6.12-9-386
drwxr-xr-x 19  root  root  4096 2005-10-14 22:30  lib
drwxr-xr-x  2  root  root 49152 2005-10-14 18:00  lost+found
drwxr-xr-x  5  root  root  4096 2005-10-14 18:01  media
drwxr-xr-x  2  root  root  4096 2005-10-05 05:37  mnt
drwxr-xr-x  2  root  root  4096 2005-10-14 18:02  opt
dr-xr-xr-x 115 root  root   0 2005-10-16 08:24  proc
drwxr-xr-x  8  root  root  4096 2005-10-16 03:47  root
drwxr-xr-x  2  root  root  4096 2005-10-14 22:29 /sbin
drwxr-xr-x  2  root  root  4096 2005-10-14 18:02  srv
drwxr-xr-x 10  root  root   0 2005-10-16 08:24  sys
drwxrwxrwt  9  root  root  4096 2005-10-16 17:27  tmp
drwxr-xr-x 13  root  root  4096 2005-10-14 22:29  usr
drwxr-xr-x 14  root  root  4096 2005-10-14 22:47  var
lrwxrwxrwx  1  root  root   25 2005-10-14 18:08  vmlinuz -> boot/vmlinuz-2.6.
12-9-386

```

The first thing you'll notice from this is that the root of the file system contains largely directories and that all files and directories are owned by root.

Only users with administrative powers can write files to the root of the file system. That means if you wanted to write to the root of the file system or otherwise access those files, you would need to use the `sudo` command. This is to prevent damage, since most of the directories in the root of the file system are vital to the correct running of Linux and contain essential programs or data.

Caution It's incredibly easy to slip up when using the command-line shell and thereby cause a lot of damage. For example, simply mistyping a forward slash in a command can mean the difference between deleting the files in a directory and deleting the directory itself. This is just another reason why you should always be careful when working at the command line, especially if you use the `sudo` command.

As you can see from the file permissions of each directory in the root of the file system, most directories allow all users to browse them and access the files within (the last three characters of the permissions read `r-x`). You just won't be able to create new files there or delete the

directories themselves. You might be able to modify or execute programs contained within the directory, but this will depend on the permissions of each individual file.

Table 14-3 provides a brief description of what each directory and file in the Ubuntu root file system contains. This is for reference only; there's no need for you to learn this information. The Ubuntu file system broadly follows the principles in the Filesystem Hierarchy Standard, as do most versions of Linux, but it does have its own subtleties.

Note The `->` symbol in Table 14-3 indicates a symbolic link, which is similar to what Windows refers to as a shortcut.

Table 14-3. *Directories and Files in the Ubuntu Root File System*

Directory	Contents
<code>bin</code>	Vital tools necessary to get the system running or for use when repairing the system and diagnosing problems
<code>boot</code>	Boot loader programs and configuration files (the boot loader is the menu that appears when you first boot Linux)
<code>cdrom -> media/cdrom</code>	Symbolic link (shortcut) to the entry for the CD or DVD-ROM drive in the <code>/dev</code> folder (accessing this file will let you access the CD-ROM drive)
<code>devbootstrap</code>	Contains files generated during the installation of Ubuntu
<code>dev</code>	Virtual files representing hardware installed on your system
<code>etc</code>	Central repository of configuration files for your system
<code>home</code>	Where each user's personal directory is stored
<code>initrd.img -> boot/initrd.img-2.6.12-9-386</code>	Symbolic link to the initial ramdisk, which is used to boot Linux
<code>lib</code>	Shared system files used by Linux as well as the software that runs on it
<code>lost+found</code>	Folder where salvaged scraps of files are saved in the event of a problematic shutdown and subsequent file system check
<code>media</code>	Where the directories representing various mounted storage systems are made available (including Windows partitions on the disk)
<code>mnt</code>	Directory in which external file systems can be temporarily mounted
<code>opt</code>	Software that is theoretically optional and not vital to the running of the system (many software packages you use daily can be found here)
<code>proc</code>	Virtual directory containing data about your system and its current status
<code>root</code>	The root user's personal directory
<code>sbin</code>	Programs essential to administration of the system

Table 14-3. *Directories and Files in the Ubuntu Root File System*

Directory	Contents
srv	Configuration files for any network servers you might have running on your system
sys	Mount point of the sysfs file system, which is used by the kernel to administer your system's hardware
tmp	Temporary files stored by the system
usr	Programs and data that might be shared with other systems (such as in a large networking setup with many users) ¹
var	Used by the system to store data that is constantly updated, such as printer spooling output
Vmlinuz -> boot/vmlinuz-2.6.12-0-386	Symbolic link to the kernel file used during bootup

¹ The `usr` directory contains its own set of directories that are full of programs and data. Many system programs, such as the X11 GUI software, are located within the `/usr` directory. Note that the `/usr` directory is used even if your system will never act as a server to other systems.

TYPES OF FILE SYSTEMS

Linux is all about choice, and this extends to the technology that makes the file system work. Unlike with Windows, where you can choose between just NTFS and FAT32 (with the emphasis being on NTFS), Linux offers many different types of file system technology. Each is designed for varying tasks. Most are scalable, however, which means that they will work just as happily on a desktop PC as on a massive cluster of computers.

Ubuntu uses the ext3 file system. This is a popular choice among distros, although the SUSE distro uses the reiserfs system. People are constantly arguing about which file system is best. The principal measuring stick is performance. Your computer spends a lot of time writing and reading files, so the faster a file system is, the faster your PC will be overall (although, in reality, the hardware is of equal importance).

Note that what we're talking here is the underlying and invisible technology of the file system. In day-to-day use, the end user won't be aware of any difference between ext3, reiserfs, or another file system technology (although when things go wrong, different tools are used to attempt repairs; their selection is automated within Ubuntu).

Here are the various types along with notes about what they offer:

- **ext2:** Fast, stable, and well established, ext2 was once the most popular type of file system technology used on Linux. It has now been eclipsed by ext3.
- **ext3:** An extension of ext2, ext3 allows journaling, a way of recording what has been written to disk so that a recovery can be attempted when things go wrong.
- **reiserfs:** This is another journaling file system, which claims to be faster than others and also offers better security features.
- **jfs:** This is a journaling file system created by IBM. It's used on industrial implementations of Unix.
- **xfs:** This is a 64-bit journaling file system created by Silicon Graphics, Inc. (SGI) and used on its own version of Unix, as well as Linux.

Mounting

Described in technical terms, mounting is the practice of making a file system available under Linux. This can take the form of a partition on your hard disk, a CD-ROM, a network server, or many other things.

Mounting drives might seem a strange concept, but it actually makes everything much simpler than it might be otherwise. For example, once a drive is mounted, you don't need to use any special commands or software to access its contents. You can use the same programs and tools that you use to access all of your other files. Mounting creates a level playing field on which everything is equal and can therefore be accessed quickly and efficiently.

Using the mount Command

Mounting is usually done via the `mount` command. Under Ubuntu, you must have administrator powers to use the `mount` command, which means prefacing it with `sudo` and providing your password.

With most modern versions of Linux, `mount` can be used in two ways: by specifying all the settings immediately after the command, or by making reference to an entry within the `fstab` file. `fstab` is a configuration file stored in the `/etc` directory that contains details of all file systems on the PC that can be mounted.

Note The root file system is itself mounted automatically during bootup, shortly after the kernel has started and has all your hardware up and running. Every file system that Linux uses must be mounted at some point.

Let's say that you insert a CD or DVD into your computer's DVD-ROM drive. To mount the CD or DVD and make it available to Linux (something that is actually done automatically as soon as you put a disk in the drive, so this example is for demonstration purposes only), you would type the following:

```
sudo mount /cdrom
```

The `mount` command first looks in your `fstab` file in the `etc` directory to find what you're referring to. Figure 14-4 shows an example of the contents of that file. (The example in the figure uses the `cat` command, which is discussed in Chapter 15). Using this information, the `mount` command attempts to make the contents of the CD available in the `/cdrom` directory. Note that this is done in a virtual way; the files are not literally copied into the directory. The directory is merely a magical conduit that allows you to read the CD's contents.

There aren't any special commands used to work with mounted drives. The shell commands discussed in Chapter 13 should do everything you need.

The `mount` command doesn't see widespread usage by most users nowadays, because most removable storage devices like CDs, and even memory card readers, are mounted automatically under Ubuntu and an icon for them appears on the desktop. However, there may be occasions when you need to mount a drive manually.

```

keir@ubuntu:~$ cat /etc/fstab
# /etc/fstab: static file system information.
#
# <file system> <mount point> <type> <options> <dump> <pass>
proc /proc proc defaults 0 0
/dev/hda2 / ext3 defaults,errors=remount-ro 0 1
/dev/hda5 none swap sw 0 0
/dev/hdc /media/cdrom0 udf,iso9660 user,noauto 0 0
/dev/fd0 /media/floppy0 auto rw,user,noauto 0 0
keir@ubuntu:~$

```

Figure 14-4. Details of all frequently mounted file systems are held in the `/etc/fstab` file.

Mounting a Drive Manually

Let's look at an example of when you might need to mount a drive manually. Suppose that you've just added a second hard disk to your PC that has previously been used on a Windows system. This has been added as the primary slave.

The first thing to do is create a mount point, which is an empty directory that will act as a location where you can tell `mount` to make the disk accessible. You can create this directory anywhere, but under Linux, the convention is to create it in the `/mnt` directory. Therefore, the following command should do the trick (note that you need to use the `sudo` command because writing to any directory other than your `/home` directory requires administrator privileges):

```
sudo mkdir /mnt/windows
```

You now need to know what kind of partition type is used on the disk, because you need to specify this when mounting. To find this out, use the `fdisk` command. Type the following exactly as it appears:

```
sudo fdisk -l /dev/hdb
```

This will list the partitions on the second disk drive (assuming an average PC system). With most hard disks used under Windows, you should find a single partition that will be either NTFS or FAT32. The examples here assume that this is `hdb1`.

Caution Be aware that `fdisk` is a dangerous system command that can damage your system. The program is designed to partition disks and can wipe out your data if you're not careful!

With this information in hand, you're now ready to mount the disk. For a FAT32 disk, type the following:

```
sudo mount -t vfat -o umask=000 /dev/hdb1 /mnt/windows
```

For an NTFS disk, type the following:

```
sudo mount -t ntfs -o umask=0222 /dev/hdb1 /mnt/windows
```

The `-t` command option is used to specify the file system type. The `-o` flag indicates that you're going to specify some more command options and you do so in the form of `umask`, which tells `mount` to ensure that the directory is readable (and writable in the case of the FAT32 drive). After this, you specify the relevant file in the `/dev` directory (this file is only virtual, of course, and merely represents the hardware), and then specify the directory that is acting as your mount point.

Note You can only read from NTFS drives under Linux because writing to them is considered too risky and might result in data loss. However, you can both read and write to FAT32 partitions because of the much simpler technology used.

Now when you browse to the `/mnt/windows` directory, by typing `cd /mnt/windows`, you should find the contents of the hard disk accessible.

For more information about the `mount` command, read its man page (type `man mount`).

Removing a Mounted System

To unmount a system, you use the special command `umount` (notice there's no `n` after the first `u`). Here's an example of using the command to unmount the CD-ROM drive:

```
sudo umount /media/cdrom
```

You need to tell the `umount` command where to find the mounted contents. You can specify either the mount directory, as in the example, or the file in the `/dev` directory that refers to mounted resource (this is a little complicated, so in most cases, it's better to simply specify the file system location of the mount).

Note Sometimes you can automatically unmount a CD-ROM by hitting the `eject` button. This is because Ubuntu is clever enough to automatically unmount the CD if the button is pressed. However, if a program is accessing the CD, or if the CD is being browsed, the button won't work until all access to the CD has ceased.

If you're currently browsing the mounted directory, you'll need to leave it before you can unmount it, no matter whether you choose to do this manually or by hitting an eject button. The same is true of all kinds of access to the mounted directory. If you're browsing the mounted drive with Nautilus or if a piece of software is accessing it, you won't be able to unmount it until you've quit the program and closed the Nautilus window (or browsed to a different part of the file system).

File Searches

Files frequently get lost. Well, technically speaking, they don't actually get lost. We just forget where we've put them. But because of this, the shell includes some handy commands to search for files.

Using the find Command

Like the Search option on the Windows XP Start menu, the `find` command manually searches through all the files on the hard disk. It's not a particularly fast way of finding a file, but it is reliable.

Here's an example:

```
find /home/keir -name "myfile"
```

This will search for `myfile` using `/home/keir` as a starting point (which is to say that it will search all directories within `/home/keir`, and any directories within those directories, and so on, because it's recursive).

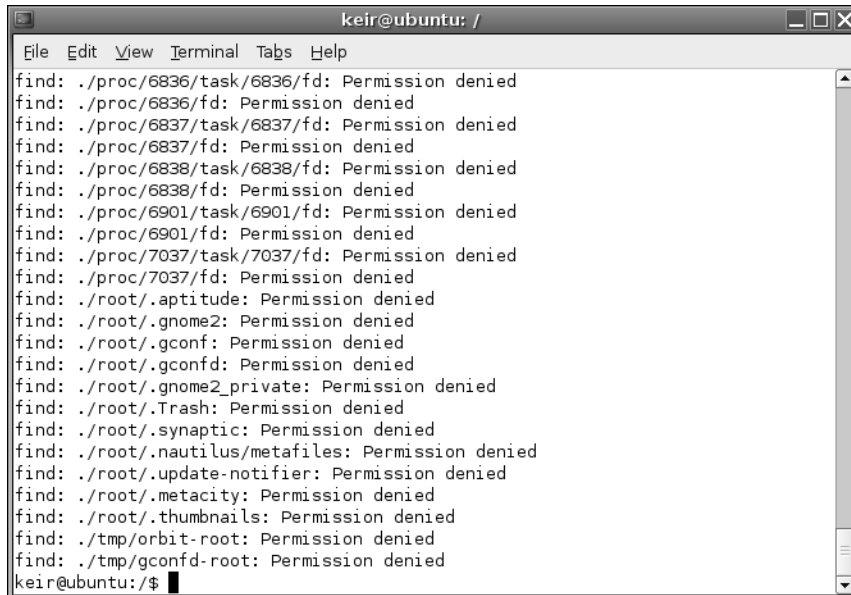
You can search the entire file system by leaving out the initial path. In this case, `find` will assume you want the search to start from `/`, the root of the file system.

If the file is found, you'll see it appear in the output of the command. The path will be shown next to the filename.

Period punctuation symbols have interesting meanings within file listings and therefore within the output of the `find` command. As you learned in Chapter 13, `..` refers to the parent directory of the one you're currently browsing. In a similar way, a single `.` refers to the directory you're in at the moment; it's shorthand for "right here." So, if `find` returns a result like `./myfile`, it means that `myfile` is right here in the current directory.

However, when a single period is used *at the beginning of a filename*, such as in `.bashrc`, it has the effect of hiding the file. In other words, that file won't appear when you type `ls` (although you can type `ls -a` to see all files, even those that are hidden).

If you give `find` a try, you'll see that it's not a particularly good way of searching. Apart from being slow, it will also return a lot of error messages about directories it cannot search. This is because, when you run the `find` command, it takes on your user permissions. Whenever `find` comes across a directory it cannot access, it will report it to you, as shown in the example in Figure 14-5. There are frequently so many of these warnings that the output can hide the instances where `find` actually locates the file in question!



```
keir@ubuntu: /
File Edit View Terminal Tabs Help
find: ./proc/6836/task/6836/fd: Permission denied
find: ./proc/6836/fd: Permission denied
find: ./proc/6837/task/6837/fd: Permission denied
find: ./proc/6837/fd: Permission denied
find: ./proc/6838/task/6838/fd: Permission denied
find: ./proc/6838/fd: Permission denied
find: ./proc/6901/task/6901/fd: Permission denied
find: ./proc/6901/fd: Permission denied
find: ./proc/7037/task/7037/fd: Permission denied
find: ./proc/7037/fd: Permission denied
find: ./root/.aptitude: Permission denied
find: ./root/.gnome2: Permission denied
find: ./root/.gconf: Permission denied
find: ./root/.gconfd: Permission denied
find: ./root/.gnome2_private: Permission denied
find: ./root/.Trash: Permission denied
find: ./root/.synaptic: Permission denied
find: ./root/.nautilus/metabytes: Permission denied
find: ./root/.update-notifier: Permission denied
find: ./root/.metacity: Permission denied
find: ./root/.thumbnails: Permission denied
find: ./tmp/orbit-root: Permission denied
find: ./tmp/gconfd-root: Permission denied
keir@ubuntu:/$
```

Figure 14-5. The `find` command is useful for finding files but isn't problem-free.

You can avoid these error messages in various ways, but perhaps the quickest solution is to preface the `find` command with `sudo` to invoke superuser powers. In this way, you'll have access to every file on the hard disk, so the `find` command will be unrestricted in where it can search and won't run into any directories it doesn't have permission to enter.

Caution Using the `sudo` command with `find` may represent an invasion of privacy if you have more than one user on your system. The `find` command will search other users' `/home` directories and report any instances of files found there, too.

However, an even better solution for finding files is to use the `locate` command.

Using the `locate` Command

The alternative to using `find` is to use the `locate` command. This is far quicker than `find` because it relies on a central database of files, which is periodically updated. In other words, it doesn't literally search the file system each time.

The problem is that if you've saved a file recently and are hoping to find it, there's a chance that it won't yet appear in `locate`'s database, so it won't turn up in the list of results.

Using `locate` is easy. You can use the following command to search for a file (you don't need to precede the command with `sudo`):

```
locate myfile
```

It's possible to update the locate database manually, although this might take a few minutes to work through. Simply issue the command:

```
sudo updatedb
```

After this, all files in the system should be indexed, making your search results more accurate.

Using the whereis Command

One other command worth mentioning in the context of searching is `whereis`. This locates where programs are stored and is an excellent way of exploring your system. Using it is simply a matter of typing something like this:

```
whereis cp
```

This will tell you where the `cp` program is located on your hard disk. It will also tell you where its source code and man page are located (if applicable). However, the first path returned by the search will be the location of the program itself.

File Size and Free Space

Often, it's necessary to know how large files are and to know how much space they're taking up on the hard disk. In addition, it's often handy to know how much free space is left on a disk.

Viewing File Sizes

Using the `ls -l` command option will tell you how large each file is in terms of bytes. Adding the `-h` option converts these file sizes to kilobytes, megabytes, and even gigabytes, depending on how large they are.

In order to get an idea of which are the largest files and which are the smallest, you can add the `-S` command option. This will order the files in the list in terms of the largest and smallest files.

The following will return a list of all the files in the current directory, in order of size (largest first), detailing the sizes in kilobytes, megabytes, or gigabytes:

```
ls -Slh
```

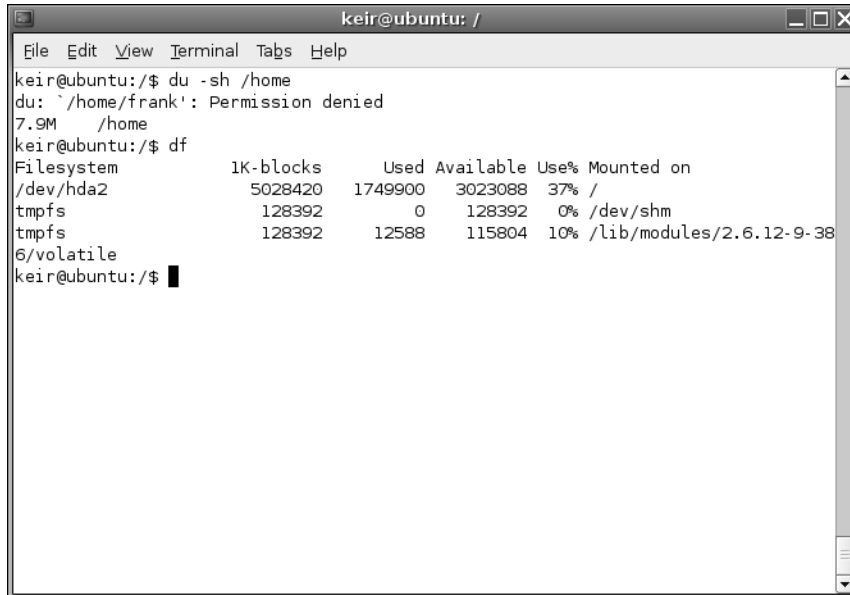
There's another more powerful way of presenting this information: using the `du` command, which stands for disk usage. When used on its own without command switches, `du` simply presents the size of directories alongside their names (starting in the current directory). It will show any hidden directories (directories whose names start with a period) and will also present a total at the end of the list. This will probably be quite a long list. Once again, you can add the `-h` command option to force the `du` command to produce human-readable measurements of kilobytes and megabytes.

If you specify a file or directory when using the `du` command, along with the `-s` command option, you can find out its total file size:

```
du -sh mydirectory
```

This will show the size taken up on the disk by `mydirectory`, adding to the total any files or subdirectories it contains.

However, `du` is limited by the same file permission problems as the `find` tool, as shown in Figure 14-6. If you run `du` as an ordinary user, it won't be able to calculate the total for any directories you don't have permission to access. Therefore, you might consider prefacing the command with `sudo`.



```

keir@ubuntu: /
File Edit View Terminal Tabs Help
keir@ubuntu:/$ du -sh /home
du: `/home/frank': Permission denied
7.9M    /home
keir@ubuntu:/$ df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/hda2              5028420    1749900   3023088   37% /
tmpfs                  128392         0    128392    0% /dev/shm
tmpfs                  128392     12588   115804   10% /lib/modules/2.6.12-9-38
6/volatile
keir@ubuntu:/$

```

Figure 14-6. The `du` command shows the size of a file, and the `df` command can be used to gauge the amount of free space on the disk.

Finding Out the Amount of Free Space

What if you want to find out how much free space is left on the disk? In this case, you can use the `df` command. This command is also demonstrated in Figure 14-6.

The `df` command works on a partition-by-partition basis. Typing it at the command prompt will show you how much space is free on the entire disk. Once again, you can add the `-h` option to the `df` command to have the file sizes returned in megabytes and gigabytes (and even terabytes if your hard disk is big enough!).

The `df` command also returns information about something called `tmpfs`. These are temporary file stores and you can ignore them.

Note There is as much space free in any directory as there is space on the disk, which is why `df` displays data about the entire partition. If you're using a system managed by a system administrator within a business environment, you might find that quotas have been used to limit how much disk space you can take up. However, if you're using a desktop PC and are the only user, this won't be activated.

Summary

In this chapter, we examined how the Ubuntu file system lies at the heart of an understanding of how the operating system works. We also discussed how the file system and user accounts go hand-in-hand and are inextricably linked. This involved discussing the concept of file ownership and usage permissions, plus how these can be manipulated using command-line shell tools.

We also discussed the overall structure of the Ubuntu file system and how external file systems can be mounted and made available within Ubuntu. Finally, we looked at how to find files and how to gauge how much free space there is within the file system.

In the next chapter, we'll look at how the BASH shell can be used to view and otherwise manipulate text files, which are also important to the way Ubuntu works.



Working with Text Files

Windows views text files as just another file type, but to Ubuntu, they can be essential components that make the system work. Configuration files are stored as plain text, and program documentation is also stored as text. This is clearly different from Windows, where it's very likely any information you're supposed to read will be contained in a Windows Help file, a rich text format (RTF) file, or even a Microsoft Word document.

Because of the reliance on text files, the shell includes several commands that let you display, edit, and otherwise manipulate text files in various ways. Learning to use the shell, and therefore learning how to administer your Ubuntu system, involves having a good understanding of these text tools. You'll use text tools for editing configuration files and viewing log files, as just two examples.

Viewing Text Files

You can easily view files using command-line tools, including `cat`, `less`, `head`, and `tail`. The simplest command for dealing with text files is `cat`.

Using the `cat` Command

When followed with a filename, the `cat` command will display the text file on screen:

```
cat mytextfile
```

`cat` is short for concatenate, and it isn't designed just to display text files. That it can do so is simply a side effect of its real purpose in life, which is to join two or more files together. However, when used with a single file, it simply displays its contents on screen.

If you try to use `cat`, you'll realize that it's good for only short text files; large files scroll off the screen.

Using the `less` Command

Because `cat` works well only with short files, and to give you more control when viewing text files, the `less` and `more` commands were created. The `more` command came first but was considered too primitive, so someone came up with `less`, which is preferred by many Linux users. However, both are usually available on the average Linux installation.

Let's look at using `less` to read the OpenOffice.Org README file, which contains information about the current release of the office suite. The file is located at `/usr/lib/openoffice2/README`, so to use `less` to read it, type the following:

```
less /usr/lib/openoffice2/README
```

You can scroll up and down within the `less` display by using the cursor keys. If you want to scroll by bigger amounts of text, you can use the Page Up and Page Down keys. Alternatively, you can use the spacebar and B key, both of which are commonly used by old-hand Linux users for the same function. In addition, the Home and End keys will take you to the start and end of the document, respectively.

When using `less`, keep an eye on the bottom part of the screen, where you'll see a brief status bar. Alongside the filename, you'll see how many lines the document has and which line you're currently up to. In addition, you'll see as a percentage the amount of document you've already read through, so you'll know how much is left to go.

`less` lets you search forward through the file by typing a slash (/), and then entering your search term. Any words that are matched will be highlighted on screen. To repeat the search, type `n`. To search backward in a file from your current point, type a question mark (?). To quit `less`, simply type `q`.

Although it's supposedly a simple program, `less` is packed with features. You can see what options are available by reading its man page or by typing `less --help`.

Using the head and tail Commands

A couple of other handy commands that you can use to view text files are `head` and `tail`. As their names suggest, these let you quickly view the beginning (head) of a file or the end (tail) of it.

Using the commands is simple:

```
tail mytextfile
```

or

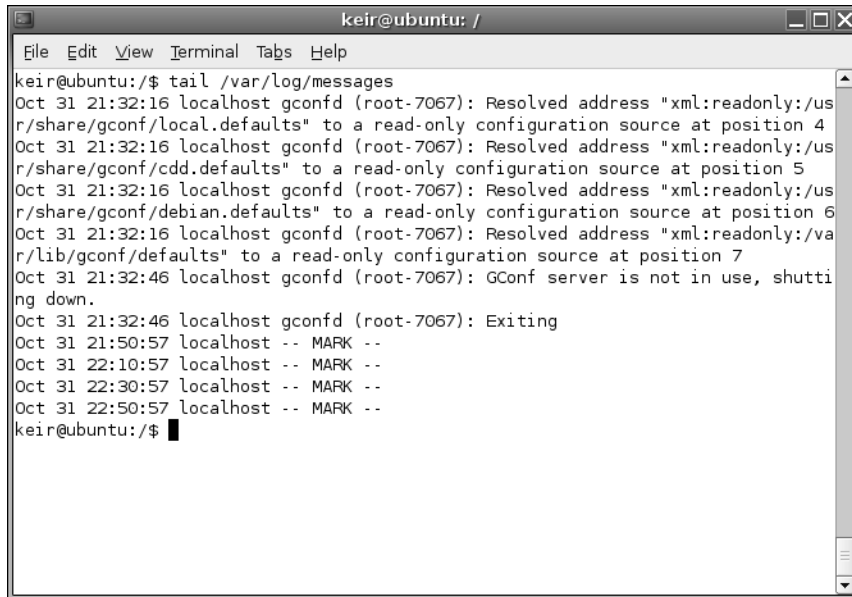
```
head mytextfile
```

By default, both commands will display ten lines of the file. You can override this by using the `-n` command option followed by the number of lines you want to see. For example, the following will show the last five lines of `mytextfile`:

```
tail -n5 mytextfile
```

These two commands are very useful when viewing log files that might contain hundreds of lines of text. The most recent information is always at the end, so `tail` can be used to see what's happened last on your system, as shown in the example in Figure 15-1.

Although they're powerful, all of these shell commands don't let you do much more than view text files. If you want to edit files, you'll need to use a text editor such as `vi`.

A terminal window titled 'keir@ubuntu: /' with a menu bar (File, Edit, View, Terminal, Tabs, Help). The terminal shows the command 'tail /var/log/messages' and its output, which consists of several lines of system logs from 'gconfd' and 'MARK' commands. The output is as follows:

```
keir@ubuntu:/$ tail /var/log/messages
Oct 31 21:32:16 localhost gconfd (root-7067): Resolved address "xml:readonly:/usr/share/gconf/local.defaults" to a read-only configuration source at position 4
Oct 31 21:32:16 localhost gconfd (root-7067): Resolved address "xml:readonly:/usr/share/gconf/cdd.defaults" to a read-only configuration source at position 5
Oct 31 21:32:16 localhost gconfd (root-7067): Resolved address "xml:readonly:/usr/share/gconf/debian.defaults" to a read-only configuration source at position 6
Oct 31 21:32:16 localhost gconfd (root-7067): Resolved address "xml:readonly:/var/lib/gconf/defaults" to a read-only configuration source at position 7
Oct 31 21:32:46 localhost gconfd (root-7067): GConf server is not in use, shutting down.
Oct 31 21:32:46 localhost gconfd (root-7067): Exiting
Oct 31 21:50:57 localhost -- MARK --
Oct 31 22:10:57 localhost -- MARK --
Oct 31 22:30:57 localhost -- MARK --
Oct 31 22:50:57 localhost -- MARK --
keir@ubuntu:/$
```

Figure 15-1. The *tail* command can be very useful for viewing the last few lines of a log file.

STANDARD INPUT AND OUTPUT

If you've read any of the Ubuntu man pages, you might have seen references to *standard input* and *standard output*. Like many things in Ubuntu, this sounds complicated but is merely a long-winded way of referring to something that is relatively simple.

Standard input is simply the device that Ubuntu normally takes input from. In other words, on the majority of desktop PCs when you're using the command-line shell, standard input refers to the keyboard. However, it's important to note that it could also feasibly refer to the mouse or any other device on your system capable of providing input; even some software can take the role of providing standard input.

Standard output is similar. It refers to the device to which output from a command is usually sent. In the majority of cases at the command line, this refers to the monitor screen, although it could feasibly be any kind of output device, such as your PC's sound card and speakers.

The man page for the *cat* command says that it will "concatenate files and print on the standard output." In other words, for the majority of desktop Ubuntu installations, it will combine (concatenate) any number of files together and print the results on screen. If you specify just one file, it will display that single file on your screen.

In addition to hardware devices, input can also come from a file containing commands, and output can also be sent to a file instead of the screen, or even sent directly to another command. This is just one reason why the command-line shell is so flexible and powerful.

Using a Command-Line Text Editor

A variety of text editors can be used within the shell, but three stand out as being ubiquitous: `ed`, `vi`, and Emacs. The first in that list, `ed`, is by far the simplest. That doesn't necessarily mean that it's simple to use or lacks powerful features, but it just doesn't match the astonishing power of both `vi` and Emacs. To call `vi` and Emacs simple text editors is to do them a disservice, because both are extremely powerful interactive environments. In particular, Emacs is considered practically an operating system in itself, and some users of Linux treat it as their shell, executing commands and performing everyday tasks, such as reading and sending e-mail from within it. There are entire books written solely about Emacs and `vi`.

Tip A fourth shell-based text editor found on many Linux systems is `nano`. This offers many word processor-like features that can be helpful if you've come to Linux from a Windows background.

The downside of all the power within Emacs and `vi` is that both packages can be difficult to learn to use. They're considered idiosyncratic by even their most ardent fans. Both involve the user learning certain unfamiliar concepts, as well as keyboard shortcuts and commands.

Although there are debates about which text editor is better and which is best, it's generally agreed that `vi` offers substantial text-editing power but isn't too all-encompassing. It's also installed by default on Ubuntu. On Ubuntu, Emacs must be installed as an optional extra. Both text editors are normally available on virtually every installation of Linux or Unix. We'll concentrate on using `vi` here.

It's important to understand that there isn't just one program called `vi`. There are many versions. The original `vi` program, supplied with Unix, is rarely used nowadays. The most common version of `vi` is a clone called `vim`, for `vi` improved, and this is the version supplied with Ubuntu. However, there are other versions, such as Elvis. Most work in a virtually identical way.

Note There's always been a constant flame war between advocates of `vi` and Emacs, as to which is better. This could be quite a vicious and desperate debate, and the text editor you used was often taken as a measure of your character! Nowadays, the battle between the two camps has softened, and the Emacs versus `vi` debate is considered an entertaining cliché of Linux and Unix use. Declaring online which text editor a user prefers is often followed by a smiley symbol to acknowledge the once-fevered emotions.

Understanding `vi` Modes

The key to understanding how `vi` works is to learn the difference between the various modes. Three modes are important: Command mode, Insert mode, and Command-Line mode.

Command Mode

Command mode is *vi*'s central mode. When the editor starts up, it's in Command mode, as shown in Figure 15-2. This lets you move around the text and delete words or lines of text. *vi* returns to Command mode after most operations. In this mode, the status bar at the bottom of the screen shows information such as the percentage progress through the document. Although you cannot insert text in this mode, you can delete and otherwise manipulate words and lines within the file. You can also move through the text using the cursor keys and the Page Up and Page Down keys.

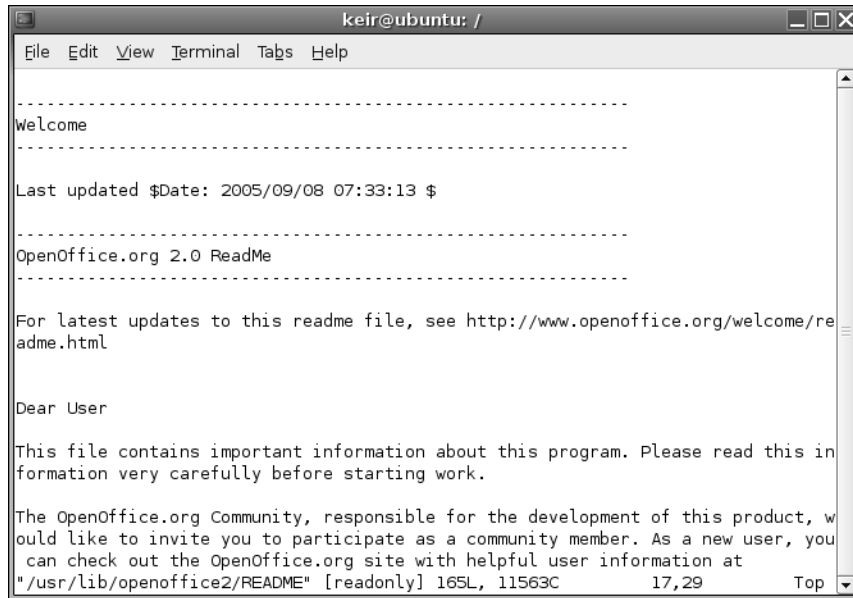


Figure 15-2. *In vi, the main mode is Command mode.*

Table 15-1 shows a list of the commands you can use in Command mode (consider photocopying it and sticking it to the side of your monitor as a handy reference).

Table 15-1. *vi Command Mode Commands*

Command	Description
Delete Text	
<i>dd</i>	Delete current line
<i>n</i> <i>dd</i>	Delete <i>n</i> number of lines (for example, <i>5dd</i> will delete five lines) ¹
<i>dw</i>	Delete the current word under the cursor ²
<i>db</i>	Delete the word before the cursor ²
<i>D</i>	Delete everything from the cursor to the end of the line ¹

Table 15-1. *vi Command Mode Commands (Continued)*

Command	Description
Search	
/	Search forward (type the search text directly after the slash)
?	Search backwards
n	Repeat search in a forward direction
N	Repeat search in a backward direction
Cut and Paste	
yy	Copy the current line ³
n yy	Copy <i>n</i> number of lines into the buffer from the cursor downwards (for example, 5yy copies five lines of text)
p	Paste the contents of the clipboard ³
Insert Text	
i	Switch to Insert mode at the cursor
o	Switch to Insert mode, placing the cursor below current line
O	Switch to Insert mode, placing the cursor above current line
A	Append text to end of line
Navigation⁴	
\$	Move the cursor to the end of the current line
w	Move the cursor to the next word
b	Move the cursor to the previous word
Miscellaneous	
.	Repeat the last command
u	Undo the last command

¹ A line ends where a line-break control character occurs in the file. Because of this, a line of text may actually take up several lines of the on-screen display.

² This will delete the remainder of current word before/after the cursor if the cursor is in the middle of a word.

³ The standard documentation refers to copying as “yanking” and the clipboard as the “buffer.”

⁴ You can also use the cursor keys to move around the file and the Page Up and Page Down keys to move up and down a page at a time. Additionally, press 0 (zero) on the main keyboard, not the numeric keypad, to move the cursor to the start of the current line, or Shift+0 to move forward one sentence (until the next full stop).

Insert Mode

To type your own text or edit text, you need to switch to Insert mode. This is normally done by typing `i`, but you can also type `O` or `o` to change to Insert mode, which is indicated by the word `INSERT` appearing at the bottom of the screen, as shown in Figure 15-3. The difference between

the commands required to switch into Insert mode is that some let you insert before or after the cursor. Generally, `i` is most useful, because what you type will appear before the character under the cursor, as with most word processors. The commands that activate Insert Mode are listed in Table 15-1, under “Insert Text.”

Tip By typing `A` (Shift+A), you can add text to the end of the line on which the cursor currently resides.

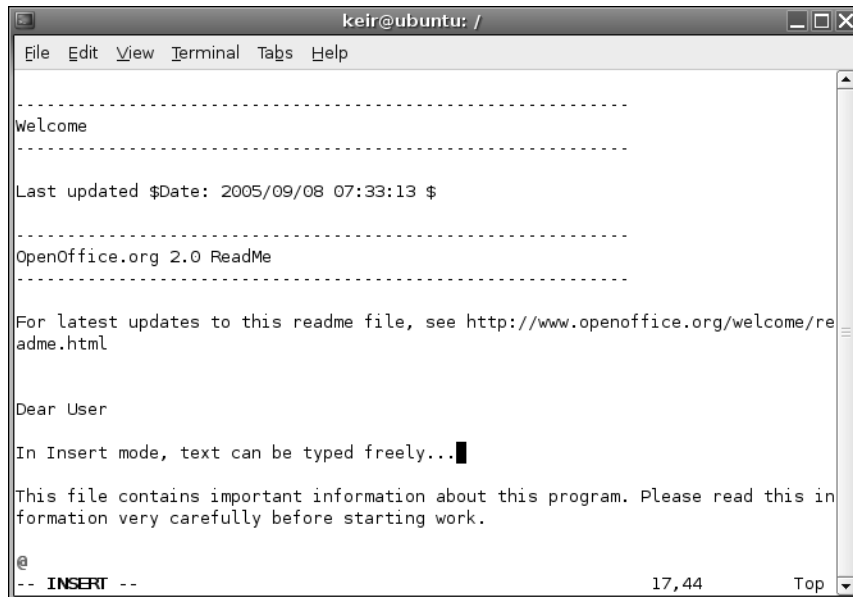


Figure 15-3. Use *vi*'s Insert mode to add and edit text.

In Insert mode, you can still move around the text using the cursor keys. Anything you type will appear at the point of the cursor. To quit this mode, press the Esc key. This will return you to Command mode.

Command-Line Mode

The third mode you should be aware of is Command-Line mode (note that, irritatingly, this is not the same as the Command mode). As its name suggests, this is the mode in which you can enter commands to save and load files, as well as perform other fundamental tasks to control *vi* or to quit the program. You can enter Command-Line mode by typing a colon (:), although if you're in Insert mode, you'll first need to leave it by pressing the Esc key. You can identify when *vi* is in this mode because the cursor will be at the bottom of the screen next to a colon symbol, as shown in Figure 15-4. To quit Command-Line mode, press the Esc key. You'll be returned to Command mode. Note that you'll automatically leave Command-Line mode after each command you issue has completed.

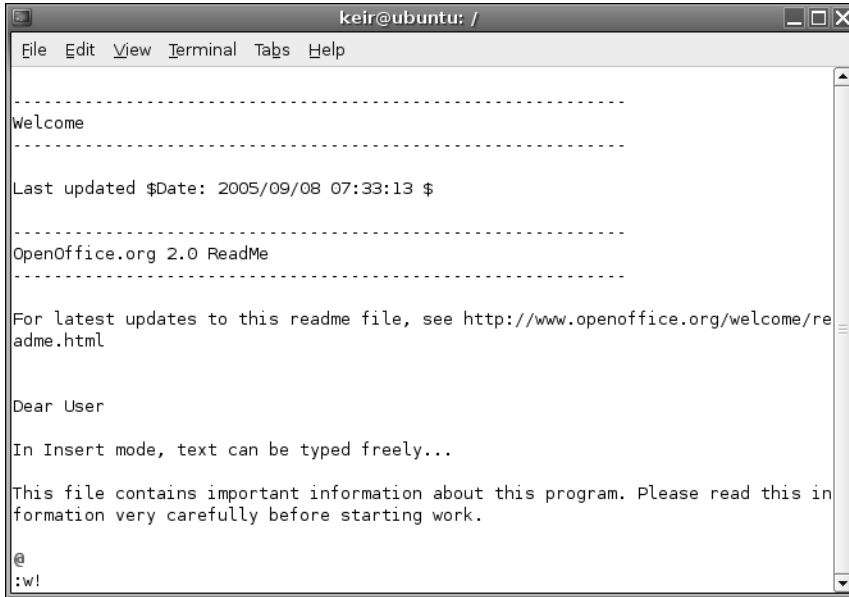


Figure 15-4. Use vi's Command-Line mode to issue commands.

For a list of basic Command-Line mode commands, see Table 15-2.

Table 15-2. Some vi Command-Line Mode Commands

Command	Description
:w	Save the file
:w!	Save the file and ignore errors such as an existing file with the same filename
:q	Quit vi
:q!	Quit vi and ignore errors such as an unsaved file
:s/word/replacement/	Search from the cursor downwards and replace any instances of the word with the replacement ¹
:help	View help documentation

¹ The search tool is very powerful and uses a number of command options for additional flexibility. Read the vi help file to learn more.

Using vi to Edit a File

As an example, let's use vi to edit the OpenOffice.org README file. You don't want to actually alter this file, so start by making a copy of it in your home directory:

```
cp /usr/lib/openoffice2/README ~
```

This will copy the file README01 to your /home directory, which you indicate using the ~ symbol.

After this, you need to change the file's permissions, because, by default, this file is read-only for all users (`r--r--r--`). Use the following command to change the permissions:

```
chmod +w README
```

Then fire up `vi` with the file, like this:

```
vi README
```

Note Windows makes a lot of use of file extensions in order to recognize files and therefore know what program to use to run them. By default, a file with a `.doc` extension tells Windows that it should use Microsoft Word to open the file, for example. Linux uses a different system based on the first few bytes of each file. Because of this, file extensions are used within Linux simply to let the users know what type of file they're dealing with. Often, they're not used at all. If a file is called `README`, you can be fairly certain that it's a text file, for example.

Once the file is opened, you'll find yourself automatically in Command mode and will be able to move around the file using the cursor keys. Altering the text is achieved using various commands (see Table 15-1). For example, typing `dd` will delete the line of text that the cursor is currently within. Typing `x` will delete the letter under the cursor. Typing `dw` will delete the current word under the cursor. Try some of these to see how they work.

To actually edit a file and type text, you'll need to switch to Insert mode. Type `i` to do this. Insert mode is fairly easy to understand. You can move around the text using the cursor keys, and then simply start typing wherever you want. The Backspace key will delete text behind the cursor, and the Delete key will delete text in front of the cursor.

When you're finished, press the Esc key to return to Command mode. Once back in Command mode, you can page through the text. The Page Up and Page Down keys will move a screenful of text at a time. Pressing the up and down cursor keys will cause the screen to scroll when the cursor reaches the top or bottom.

After you're finished editing, you'll need to save the file. This is done in Command-Line mode. You can enter this mode by typing a colon (`:`). You'll see a colon appear at the bottom of the screen, and this is where you type the commands. Note that after you type a command, you'll immediately exit Command-Line mode, so if you want to issue another command, you'll need to type a colon again.

To save a file, in Command-Line mode, type `:w` (which stands for "write"). If you want to save the current file with a different name, you'll need to enter a filename after the `w` command, like this:

```
:w mytextfile
```

To quit `vi`, type `:q`. However, if you've edited a file, you won't be able to quit until the file has been saved. If you want to save the file and then quit, you can type `:wq`. If you don't want to save the file, type `:q!`. The exclamation point tells `vi` to override any objections it might have. You can also use it with the save command—`:w!`—to force the overwriting of a file that already exists.

Note If you don't have the correct permissions to write a file, `vi` might tell you that you can use `:w!` to override. In this case, it's wrong. The only way to write to a file for which you don't have permissions is to change its permissions.

Creating a New Text File Using `vi`

Creating and editing a new file with `vi` is easy. From any command-line shell, simply type this:

```
vi myfile
```

This will start `vi` and give your new file a name. However, the file won't be saved until you manually issue the save command (`:w`) in `vi`. This means that if your computer crashes before you save, the file will be lost!

Note The version of `vi` provided with Ubuntu, `vim`, includes some elementary file-save protection. If, for any reason, `vim` is not shut down properly, there's a chance you'll be able to recover a version of file the next time `vim` starts. However, as with all such protection in any kind of program, you shouldn't rely on this. You should use the `:w` command to save your file periodically.

As always with `vi`, you start out in the default Command mode. To start typing immediately, enter Insert mode by typing `i`. You'll notice when typing that although the text is wrapped on each line, words are not carried over, and they often break across lines in an ugly way. This is because `vi` is primarily a text editor, not a word processor. For people who create text files, like programmers, having line breaks shown in this way can be useful.

When you're finished typing a sentence or paragraph, you can press the Enter key as usual to start a new line. You should then be able to move between lines using the up and down cursor keys. You'll notice an odd thing when you try to do this, however: unlike with a word processor, moving up a line of text that spreads across more than one line on screen will take the cursor to the start of the line, rather than into the middle of it. This again relates to `vi`'s text editor focus, where such a feature is useful when editing documents such as program configuration files.

When you're finished, press the Esc key to switch to Command mode. Then type a colon to enter Command-Line mode. Type `:w` to save the file using the filename you gave it earlier. If you started `vi` without specifying a filename, you'll need to specify a filename with the save command, such as `:w myfile`.

USING GEDIT TO EDIT TEXT FILES

If all this talk of `vi` sounds like too much hard work, don't forget that the GNOME desktop includes an excellent text editor in the form of Gedit. In fact, to describe Gedit as merely a text editor is to do it something of a disservice because it includes many handy word processor-like features.

You can call Gedit and open a file in it from the command-line prompt as follows:

```
gedit <filename>
```

If you need to adopt superuser powers to edit the likes of configuration files, simply preface it with `sudo`:

```
sudo gedit <filename>
```

You'll find Gedit fairly straightforward to use.

Searching Through Files

You can search for particular words or phrases in text files by loading the file into `less` or `vi` (see Table 15-1). The maneuverability offered by both programs lets you leap from point to point in the text, and their use is generally user-friendly.

However, using `vi` or `less` can take precious seconds. There's a quicker command-line option that will search through a file in double-quick speed: `grep`.

Using `grep` to Find Text

`grep` stands for Global Regular Expression Print. `grep` is an extremely powerful tool that can use pattern-based searching techniques to find text in files. Pattern-based searching means that `grep` offers various options to loosen the search so that more results are returned.

The simplest way of using `grep` is to specify some brief text, followed by the name of the file you want to search. Here's an example:

```
grep 'helloworld' myfile
```

This will search for the phrase `helloworld` within `myfile`. If it's found, the entire line that `helloworld` is on will be displayed on screen.

If you specify the `*` wildcard instead of a filename, `grep` will search every file in the directory for the text. Adding the `-r` command option will cause `grep` to search all the files, and also search through any directories that are present:

```
grep -r 'helloworld' *
```

Another handy command option is `-i`, which tells `grep` to ignore uppercase and lowercase letters when it's searching. Figure 15-5 shows an example of using `grep`.

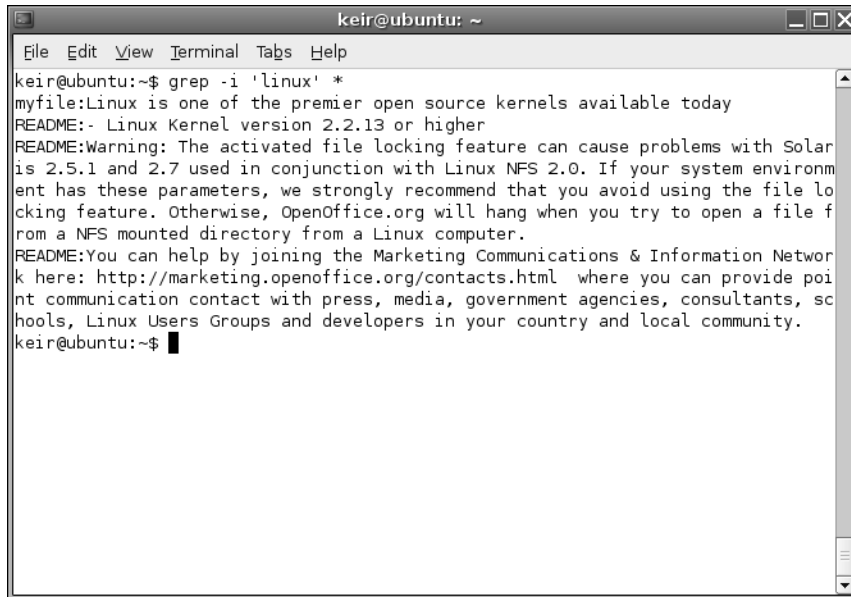
A screenshot of a terminal window titled "keir@ubuntu: ~". The window has a menu bar with "File", "Edit", "View", "Terminal", "Tabs", and "Help". The terminal shows the command `grep -i 'linux' *` and its output. The output includes the text "myfile:Linux is one of the premier open source kernels available today", followed by a README for Linux Kernel version 2.2.13 or higher. The README contains a warning about file locking features in Solaris 2.5.1 and 2.7, and a link to the Marketing Communications & Information Network at <http://marketing.openoffice.org/contacts.html>. The terminal ends with the prompt `keir@ubuntu:~$` and a cursor.

Figure 15-5. *grep* is a powerful tool that can search for text within files.

Tip You might never choose to use `grep` for searching for text within files, but it can prove very handy when used to search through the output of other commands. This is done by “piping” the output from one command to another, as explained in Chapter 17.

Using Regular Expressions

The true power of `grep` is achieved by the use of search patterns known as *regular expressions*, or *regexes* for short. Put simply, regexes allow you to be vague rather than specific when searching, meaning that `grep` (and many similar tools that use the system of regexes, such as the `find` command discussed in Chapter 14) will return more results.

For example, you can specify a selection or series of characters (called a *string* in regex terminology) that might appear in a word or phrase you’re searching for. This can be useful if you’re looking for a word that might be spelled differently from how you anticipate, for example.

The most basic form of regex is the bracket expansion. This is where additional search terms are enclosed in square brackets within a search string. For example, suppose you want to find a file that refers to several drafts of a document you’ve been working on. The files are called `myfile_1draft.doc`, `myfile_2draft.doc`, and so on. To find any document that mentions these files, you could type this:

```
grep 'myfile_[1-9]draft\.doc' *
```

The use of square brackets tells `grep` to fill in details within the search string based on what’s inside the square brackets. In this case, `1-9` means that all the numbers from one to nine should be applied to the search string. It’s as if you’ve told `grep` to search for `myfile_1draft.doc`, and

then told it to search for `myfile_2draft.doc`, and so on. Notice that the example has a backslash before the period separating the file extension from the filename. This indicates to `grep` that it should interpret the period as an element of the string to be searched for, rather than as a wildcard character, which is how `grep` usually interprets periods.

You don't need to specify a range of characters in this way. You can simply enter whatever selection of characters you want to substitute into the search string. Here's an example:

```
grep 'myfile[12345]\.doc' *
```

This will attempt to find any mention of `myfile1.doc`, `myfile2.doc`, `myfile3.doc`, and so on, in any file within the directory.

Here's another example:

```
grep '[GgNn]ome' *
```

This will let you search for the word `Gnome` within files but takes into account any possible misspelling of the word by people who forget to use the silent `G`, and any use of uppercase or lowercase.

This is only scratching the surface of what regexes can do. For example, many regexes can be combined together into one long search string, which can provide astonishing accuracy when searching. Table 15-3 contains some simple examples that should give you an idea of the power and flexibility of regexes.

Table 15-3. *Some Examples of Regular Expressions*

Search String	Description
'document[a-z]'	Returns any lines containing the string “document” followed by any single letter from the range <i>a</i> through <i>z</i> .
'document[A-Za-z]'	Returns any lines containing the string “document” followed by the letters <i>A</i> through <i>Z</i> or <i>a</i> through <i>z</i> . Note that no comma or other character is needed to separate possibilities within square brackets.
'document.'	Returns any lines containing the string “document” followed by any other character. The period is used as a wildcard signifying any single character.
'document[[:digit:]]'	Returns any lines containing the string “document” followed by any number.
'document[[:alpha:]]'	Returns any lines containing the string “document” followed by any character.
'^document'	Returns any lines that have the string “document” at the beginning. The caret symbol (^) tells <code>grep</code> to look only at the beginning of each line.
'document\$'	Returns any line that has the string “document” at the end of the line. The dollar sign symbol (\$) tells <code>grep</code> to look for the string only at the end of lines.
'document[^1-6]'	Returns lines that have the string “document” in them but not if it's followed by the numbers 1 through 6. When used in square brackets, the caret character (^) produces a nonmatching list—a list of results that don't contain the string.

`grep` is very powerful. It can be complicated to master, but it offers a lot of scope for performing extremely precise searches that ensure you find only what you're seeking. It's well worth reading through its man pages. You can also refer to books on the subject, of which there are many. A good example is *Regular Expression Recipes: A Problem-Solution Approach*, by Nathan A. Good (1-59059-441-X; Apress, 2004).

Comparing Text Files

If you want to compare the differences between two text files, one way to do this is to use the `diff` command. This is designed primarily to uncover small changes in otherwise identical documents, such as revisions made by another person. Of course, it can also be used to prove that two files are identical. If you run the files through `diff`, and it shows no output, it has been unable to spot any differences.

`diff` is ordinarily used like this:

```
diff mytextfile1 mytextfile2
```

If `diff` spots any differences between the files, the results are a little more complicated than you might be used to. Any lines that are different within the files will appear on screen. Those lines that are identical won't be displayed. Lines preceded with a left angle bracket (<) are from the first file, while those with a right angle bracket (>) are from the second file.

For a different display, you could type something like this:

```
diff -y mytextfile1 mytextfile2
```

This places the two lists side-by-side and highlights lines that are different with a pipe symbol (|). However, it requires a lot more screen space than using `diff` without the `-y` option.

Note When you use the `-y` command option with `diff`, it will struggle to fit the output in a standard GNOME Terminal window. If it is maximized on a 17-inch screen (1024×768 resolution), it should be just large enough to fit the information in, depending on the complexity of the files being compared.

By specifying the `-a` command option, you can make `diff` process binary files, too. This is a handy way of comparing virtually any kind of files, including program files, to see if they're identical. If there's no output from `diff`, then the two files are identical. If your screen fills with gibberish, then the files are clearly different.

Incidentally, if you want to compare three documents, you can use a very similar command: `diff3`. Check the command's man page to learn more about how it works.

Summary

In this chapter, we examined how text files can be manipulated. In many ways, the BASH shell is built around manipulating text, and we explored various tools created with this goal in mind. We started with the commands that can display text files (or part of them).

We then looked at how the `vi` text editor can be used to both edit and create documents. Next, we explored how regexes can be used with the `grep` command to create sophisticated search strings, which can uncover any text within documents. Finally, you saw how to compare text files.

In the next chapter, we'll look at how you can use various command-line tools to take control of your system.



Taking Control of the System

By now, you should be starting to realize that the shell offers an enormous amount of power when it comes to administering your PC. The BASH shell commands give you quick and efficient control over most aspects of your Linux setup. However, the shell truly excels in one area: controlling the processes on your system.

Controlling processes is essential for administration of your system. You can tidy up crashed programs, for example, or even alter the priority of a program so that it runs with a little more consideration for other programs. Unlike with Windows, this degree of control is not considered out of bounds. This is just one more example of how Linux provides complete access to its inner workings and puts you in control.

Without further ado, let's take a look at what can be done.

Viewing Processes

A process is something that exists entirely behind the scenes. When the user runs a program, one or many processes might be started, but they're usually invisible unless the user specifically chooses to manipulate them. You might say that programs exist in the world of the user, but processes belong in the world of the system.

Processes can be started not only by the user, but also by the system itself to undertake tasks such as system maintenance, or even to provide basic functionality, such as the GUI system. Many processes are started when the computer boots up, and then they sit in the background, waiting until they're needed (such as programs that send mail). Other processes are designed to work periodically to accomplish certain tasks, such as ensuring system files are up-to-date.

You can see what processes are currently running on your computer by running the `top` program. Running `top` is simply a matter of typing the command at the shell prompt.

As you can see in Figure 16-1, `top` provides very comprehensive information and can be a bit overwhelming at first sight. However, the main area of interest is the list of processes (which `top` refers to as *tasks*).

```

keir@ubuntu: ~
File Edit View Terminal Tabs Help
top - 02:58:00 up 53 min, 2 users, load average: 0.08, 0.66, 0.68
Tasks: 70 total, 1 running, 69 sleeping, 0 stopped, 0 zombie
Cpu(s): 10.4% us, 7.4% sy, 0.0% ni, 78.9% id, 0.3% wa, 3.0% hi, 0.0% si
Mem: 256784k total, 238692k used, 18092k free, 43832k buffers
Swap: 257000k total, 6248k used, 250752k free, 59368k cached

  PID USER      PR  NI  VIRT  RES  SHR  S  %CPU  %MEM   TIME+  COMMAND
 6437 root        15   0 41072  18m 4476  S  13.7   7.4  14:06.43 Xorg
 6928 keir        15   0 13276  7356 5836  S   1.0   2.9   0:00.64 metacity
 6935 keir        15   0 29880  14m 10m   S   1.0   6.0   0:00.93 nautilus
 6978 keir        15   0 30996  12m 8040  S   1.0   5.0   0:07.54 gnome-terminal
 6889 keir        15   0 19344  9004 6580  S   0.3   3.5   0:00.67 gnome-settings-
 6933 keir        15   0 20552  11m 8660  S   0.3   4.6   0:00.65 gnome-panel
 6937 keir        15   0 17096  7028 5436  S   0.3   2.7   0:00.13 gnome-volume-ma
 6943 keir        15   0 18628  9780 7220  S   0.3   3.8   0:00.47 wnck-applet
 6975 keir        15   0 21592  9616 7332  S   0.3   3.7   0:00.45 clock-applet
    1 root        16   0  1564   532  460  S   0.0   0.2   0:00.85 init
    2 root        34  19     0     0     0  S   0.0   0.0   0:00.00 ksoftirqd/0
    3 root        10  -5     0     0     0  S   0.0   0.0   0:00.37 events/0
    4 root        11  -5     0     0     0  S   0.0   0.0   0:00.05 khelper
    5 root        14  -5     0     0     0  S   0.0   0.0   0:00.00 kthread
    7 root        20  -5     0     0     0  S   0.0   0.0   0:00.00 kacpid
   63 root        10  -5     0     0     0  S   0.0   0.0   0:01.21 kblockd/0
   91 root        15   0     0     0     0  S   0.0   0.0   0:00.00 pdfflush

```

Figure 16-1. The *top* program gives you an eagle-eye view of the processes running on your system.

Here's an example of a line taken from *top* on my test PC, shown with the column headings from the process list:

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
5499	root	15	0	78052	25m	60m	S	2.3	5.0	6:11.72	Xorg

A lot of information is presented here, as described in Table 16-1.

Table 16-1. The *top* Program Process Information

Column	Description
PID	The first number is the process ID (PID). This is the unique number that the system uses to track the process. The PID comes in handy if you want to kill (terminate) the process (as explained in the next section of this chapter).
USER	This column lists the owner of the particular process. As with files, all processes must have an owner. A lot of processes will appear to be owned by the root user. Some of them are system processes that need to access the system hardware, which is something only the root user is allowed to do. Other processes are owned by root for protection; root ownership means that ordinary users cannot tamper with these processes.
PR	This column shows the priority of the process. This is a dynamic number, showing where the particular process is in the CPU queue at the present time.

Table 16-1. *The top Program Process Information*

Column	Description
NI	This column shows the “nice” value of the process. This refers to how charitable a process is in its desire for CPU time. A high figure here (up to 19) indicates that the process is willing to be interrupted for the sake of other processes. A negative value means the opposite: the process is more aggressive than others in its desire for CPU time. Some programs need to operate in this way, and this is not necessarily a bad thing.
VIRT	This column shows the amount of virtual memory used by the process. ¹
RES	This column shows the total amount of physical memory used. ¹
SHR	This column shows the amount of shared memory used. This refers to memory that contains code that is relied on by other processes and programs.
S	This column shows the current status of the task. Generally, the status will either be sleeping, in which case an S will appear, or running, in which case an R will appear. Most processes will be sleeping, even ones that appear to be active. Don’t worry about this; it just reflects the way the Linux kernel works. A Z in this column indicates a zombie process (a child of a process that has been terminated).
%CPU	This column shows the CPU use, expressed as a percentage. ²
%MEM	This column shows the memory use, again expressed as a percentage. ²
TIME+	This column shows a measure of how long the process has been up and running.
COMMAND	This shows the actual name of the process itself.

¹ Both VIRT and RES are measured in kilobytes unless an m appears alongside the number; in which case, you should read the figure as megabytes.

² The %CPU and %MEM entries tell you in easy-to-understand terms how much of the system resources a process is taking up.

This list will probably be longer than the screen has space to display, so `top` orders the list of processes by the amount of CPU time the processes are using. Every few seconds, it updates the list. You can test this quite easily. Let your PC rest for a few seconds, without touching the mouse or typing. Then move the mouse around for a few seconds. You’ll see that the process called Xorg leaps to the top of the list (or appears very near the top). Xorg is the program that provides the graphical subsystem for Linux, and making the mouse cursor appear to move around the screen requires CPU time. When nothing else is going on, moving the mouse causes Xorg to appear as the number one user of CPU time on your system.

Tip Typing `d` while `top` is running lets you alter the update interval, which is the time between screen updates. The default is three seconds, but you can reduce that to one second or even less if you wish. However, a constantly updating `top` program starts to consume system resources and can therefore skew the diagnostic results you’re investigating. Because of this, a longer, rather than shorter, interval is preferable.

It's possible to alter the ordering of the process list according to other criteria. For example, you can list the processes by the quantity of memory they're using, by typing `M` while `top` is up and running. You can switch back to CPU ordering by typing `P`.

RENICING A PROCESS

You can set how much CPU time a process receives while it's actually running. This is done by *renicing* the process. This isn't something you should do on a regular basis, but it can prove very handy if you start a program that then uses a lot of system resources and makes the system unbearably slow.

The first thing to do is to use `top` to spot the process that needs to be restrained and find out its PID number. This will be listed on the left of the program's entry on the list. Once you know this, type `r`, and then type in the PID number. You'll then be asked to specify a renice value. The scale goes from `-20`, which is considered the highest priority, to `19`, which is considered the lowest. Therefore, you should type `19`. After this, you should find some responsiveness has returned to the system, although how much (if any) depends on the nature of the programs you're running.

You might be tempted to bump up the priority of a process to make it run faster, but this may not work because of complexities in the Linux kernel. In fact, it might cause serious problems. Therefore, you should renice with care and only when you must.

Controlling Processes

Despite the fact that processes running on your computer are usually hidden away, Linux offers complete, unrestricted, and unapologetic control over them. You can terminate processes, change their properties, and learn every item of information there is to know about them.

This provides ample scope for damaging the currently running system but, in spite of this, even standard users have complete control over processes that they personally started (one exception is zombie processes, described a bit later in this section). As you might expect, the root user (or any user who adopts superuser powers) has control over all processes that were created by ordinary users, as well as those processes started by the system itself.

The user is given this degree of control over processes in order to enact repairs when something goes wrong, such as when a program crashes and won't terminate cleanly. It's impossible for standard users to damage the currently running system by undertaking such work, although they can cause themselves a number of problems.

Note This control over processes is what makes Linux so reliable. Because any user can delve into the workings of the kernel and terminate individual processes, crashed programs can be cleaned up with negligible impact on the rest of the system.

Killing Processes

Whenever you quit a program or, in some cases, when it completes the task you've asked of it, it will terminate itself. This means ending its own process and also that of any other processes

it created in order to run. The main process is called the *parent*, and the ones it creates are referred to as *child* processes.

Tip You can see a nice graphical display of which parent owns which child process by typing `ps tree` at the command-line shell.

While this should mean your system runs smoothly, badly behaved programs sometimes don't go away. They stick around in the process list. Alternatively, you might find that a program crashes and so isn't able to terminate itself. In very rare cases, some programs that appear otherwise healthy might get carried away and start consuming a lot of system resources. You can tell when this happens because your system will start slowing down for no reason, as less and less memory and/or CPU time is available to run actual programs.

In all of these cases, the user usually must kill the process in order to terminate it manually. This is easily done using `top`.

The first task is to track down the crashed or otherwise problematic process. In `top`, look for a process that matches the name of the program, as shown in Figure 16-2. For example, the Firefox web browser generally runs as a process called `firefox-bin`.

```
MEM  TIME+  COMMAND
1.0  14:09.19  Xorg
1.7  0:01.16  firefox-bin
1.9  0:00.73  metacity
1.8  0:00.54  wnck-applet
1.0  0:07.73  gnome-terminal
1.5  0:00.70  gnome-settings-
1.0  0:01.22  kblockd/0
1.0  0:00.12  kswapd0
1.7  0:00.71  gnome-panel
1.0  0:00.96  nautilus
1.7  0:00.33  update-notifier
```

Figure 16-2. You can normally identify a program by its name in the process list.

Caution You should be absolutely sure that you know the correct process before killing it. If you get it wrong, you could cause other programs to stop running.

Because `top` doesn't show every single process on its screen, tracking down the trouble-causing process can be difficult. A handy tip is to make `top` show only the processes created by the user you're logged in under. This will remove the background processes started by root. You can do this within `top` by typing `u`, and then entering your username.

Once you've spotted the crashed process, make a note of its PID number, which will be at the very left of its entry in the list. Then type `k`. You'll be asked to enter the PID number. Enter that number, and then press Enter once again (this will accept the default signal value of 15, which will tell the program to terminate).

With any luck, the process (and the program in question) will disappear. If it doesn't, the process you've killed might be the child of another process that also must be killed. To track down the parent process, you need to configure `top` to add the PPID field, for the parent process ID, to its display. To add this field, type `f`, and then `b`. Press Enter to return to the process list. The PPID column will appear next to the process name on the right of the window. It simply shows the PID of the parent process. You can use this information to look for the parent process within the main list of processes.

The trick here is to make sure that the parent process isn't something that's vital to the running of the system. If it isn't, you can safely kill it. This should have the result of killing the child process you uncovered prior to this.

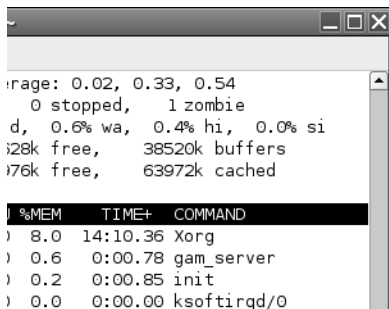
Caution If the PPID field in `top` displays a value of 1, that means the process doesn't have a parent process. In both the PPID and PID fields, you should always watch out for low numbers, particularly one-, two- or three-digit numbers. These are usually processes that started early on when Linux booted and that are essential to the system.

Controlling Zombie Processes

Zombie processes are those that are children of processes that have terminated. However, for some reason, they failed to take their child processes with them. Zombie processes are rare on most Linux systems.

Despite their name, zombie processes are harmless. They're not actually running and don't take up system resources. However, if you want your system to be spick-and-span, you can attempt to kill them.

In the top-right area of `top`, you can see a display that shows how many zombie processes are running on your system, as shown in Figure 16-3. Zombie processes are easily identified because they have a Z in the status (S) column within `top`'s process list. To kill a zombie process, type `k`, and then type its PID. Then type `9`, rather than accept the default signal of 15.



```

rverage: 0.02, 0.33, 0.54
  0 stopped,  1 zombie
d, 0.6% wa,  0.4% hi,  0.0% si
128k free,  38520k buffers
176k free,  63972k cached

%MEM  TIME+  COMMAND
)  8.0  14:10.36 Xorg
)  0.6  0:00.78 gam_server
)  0.2  0:00.85 init
)  0.0  0:00.00 ksoftirqd/0
  
```

Figure 16-3. You can see at a glance how many zombie processes are on your system by looking at the top right of `top`'s display.

Note No magic is involved in killing processes. All that happens is that `top` sends them a “terminate” signal. In other words, it contacts them and asks them to terminate. By default, all processes are designed to listen for commands such as this; it’s part and parcel of how programs work under Linux. When a program is described as *crashed*, it means that the user is unable to use the program itself to issue the terminate command (such as `Quit`). A crashed program might not be taking input, but its *processes* will probably still be running.

In many cases, zombie processes simply won’t go away. When this happens, you have two options. The first is to restart the program that is likely to be the zombie’s owner, in the hope that it will reattach with the zombie, and then quit the program. With any luck, it will take the zombie child with it this time. Alternatively, you can simply reboot your PC. But it’s important to note that zombie processes are harmless and can be left in peace on your system!

Using Other Commands to Control Processes

You don’t always need to use `top` to control processes. A range of quick and cheerful shell commands can diagnose and treat process problems.

The first of these is the `ps` command. This stands for Process Status and will report a list of currently running processes on your system. This command is normally used with the `-aux` options:

```
ps -aux
```

This will return a list something like what you see when you run `top`.

If you can spot the problematic process, look for its PID and issue the following command:

```
kill <PID number>
```

For example, to kill a process with a PID of 5122, you would type this:

```
kill 5122
```

If, after this, you find the process isn’t killed, then you should use the `top` program, as described in the previous sections, because it allows for a more in-depth investigation.

Another handy process-killing command lets you use the actual process name. The `killall` command is handy if you already know from past experience what a program’s process is called. For example, to kill the process called `firefox-bin`, which is the chief process of the Firefox web browser, you would use the following command:

```
killall firefox-bin
```

Caution Make sure you’re as specific as possible when using the `killall` command. Issuing a command like `killall bin` will kill all processes that might have the word `bin` in them!

CLEARING UP CRASHES

Sometimes, a crashed process can cause all kinds of problems. The shell you're working at may stop working, or the GUI itself might stop working properly.

In cases like this, it's important to remember that you can have more than one instance of the command-line shell up and running at any one time. For example, if a process crashes and locks up GNOME Terminal, simply start a new instance of GNOME Terminal (Applications ► Accessories ► Terminal). Then use `top` within the new window to kill the process that is causing trouble for the other terminal window.

If the crashed program affects the entire GUI, you can switch to a virtual console by pressing `Ctrl+Alt+F1`. Although the GUI disappears, you will not have killed it, and no programs will stop running. Instead, you've simply moved the GUI to the background while a shell console takes over the screen. Then you can use the virtual console to run `top` and attempt to kill the process that is causing all the problems. When you're ready, you can switch back to the GUI by pressing `Ctrl+Alt+F7`.

Controlling Jobs

Whenever you start a program at the shell, it's assigned a job number. *Jobs* are quite separate from processes and are designed primarily for users to understand what programs are running on the system.

You can see which jobs are running at any one time by typing the following at the shell prompt:

```
jobs
```

When you run a program, it usually takes over the shell in some way and stops you from doing anything until it's finished what it's doing. However, it doesn't have to be this way. Adding an ampersand symbol (&) after the command will cause it to run in the background. This is not much use for commands that require user input, such as `vi` or `top`, but it can be very handy for commands that churn away until they're completed.

For example, suppose that you want to decompress a large zip file. For this, you can use the `unzip` command. As with Windows, decompressing large zip files can take a lot of time, during which time the shell would effectively be unusable. However, you can type the following to retain use of the shell:

```
unzip myfile.zip &
```

When you do this, you'll see something similar to the following, although the four-digit number will be different:

```
[1] 7483
```

This tells you that `unzip` is running in the background and has been given job number 1. It also has been given process number 7483 (although bear in mind that when some programs start, they instantly kick off other processes and terminate the one they're currently running, so this won't necessarily be accurate).

Tip If you've ever tried to run a GUI program from the shell, you might have realized that the shell is inaccessible while it's running. Once you quit the GUI program, the control of the shell will be returned to you. By specifying that the program should run in the background with the & (ampersand symbol), you can run the GUI program and still be able to type away and run other commands.

You can send several jobs to the background, and each one will be given a different job number. In this case, when you wish to switch into a running job, you can type its number. For example, the following command will switch you to the background job assigned the number 3:

```
%3
```

You can exit a job that is currently running by pressing Ctrl+Z. It will still be there in the background, but it won't be running (officially, it's said to be *sleeping*). To restart it, you can switch back to it, as just described. Alternatively, you can restart it but still keep it in the background. For example, to restart job 2 in the background, leaving the shell prompt free for you to enter other commands, type the following:

```
%2 &
```

You can bring the command in the background into the foreground by typing the following:

```
fg
```

When a background job has finished, something like the following will appear at the shell:

```
[1]+  Done                unzip myfile.zip
```

Using jobs within the shell can be a good way of managing your workload. For example, you can move programs into the background temporarily while you get on with something else. If you're editing a file in *vi*, you can press Ctrl+Z to stop the program. It will remain in the background, and you'll be returned to the shell, where you can type other commands. You can then resume *vi* later on by typing *fg* or typing % followed by its job number.

Summary

This chapter has covered taking complete control of your system. We've looked at what processes are, how they're separate from programs, and how they can be controlled or viewed using programs such as *top* and *ps*. In addition, we explored job management under BASH. You saw that you can stop, start, and pause programs at your convenience.

In the next chapter, we'll take a look at several tricks and techniques that you can use with the BASH shell to finely hone your command-line skills.



Cool Shell Tricks

The BASH shell is the product of many years of development work by a lot of people. It comes from the old days of Unix and was an important step in computer software evolution. It's a program that retains complete simplicity, yet packs in more features than most users could ever hope to use.

One of the best things about the shell is its sheer power. If you ever wonder if you can do a task differently (and more efficiently), you'll probably find that one of the many BASH developers has implemented a method to do so. Once you learn these techniques, you'll find you can whiz around the shell at blinding speed. It's just a matter of exploring the far reaches of the shell, and that's what you'll do in this chapter. Hold onto your hats, because it's an exciting ride!

Using Autocompletion

The Tab key is your best friend when using the shell, because it will cause BASH to automatically complete whatever you type. For example, if you want to run Ubuntu's web browser, you can enter `firefox` at the command line. However, to save yourself some time, you can type `fir`, and then press Tab. You'll then find that BASH fills in the rest for you. It does this by caching the names of the programs you might run according to the directories listed in your `$PATH` variable (see Chapter 13).

Of course, autocompletion has some limitations. On my Ubuntu test system, typing `loc` didn't autocomplete `locate`. Instead, it caused BASH to beep. This is because on a default Ubuntu installation, there is more than one possible match. Pressing Tab again shows those matches. Depending on how much you type (how much of an initial clue you give BASH), you might find there are many possible matches.

In this case, the experienced BASH user simply types another letter, which will be enough to distinguish the almost-typed word from the rest, and presses Tab again. With any luck, this should be enough for BASH to fill in the rest.

Autocompletion with Files and Paths

Tab autocompletion also works with files and paths. If you type the first few letters of a folder name, BASH will try to fill in the rest. This also obviously has limitations. There's no point in typing `cd myfol` and pressing Tab if there's nothing in the current directory that starts with the letters `myfol`. This particular autocomplete function works by looking at your current directory and seeing what's available.

Alternatively, you can specify an initial path for BASH to use in order to autocomplete. Typing `cd /ho` and pressing Tab will cause BASH to autocomplete the path by looking in the root directory (`/`). In other words, it will autocomplete the command with the directory `home`. In a similar way, typing `cd myfolder/myfo` will cause BASH to attempt to autocomplete by looking for a match in `myfolder`.

If you want to run a program that resides in the current directory, such as one you've just downloaded for example, typing `./`, followed by the first part of the program name, and then pressing Tab should be enough to have BASH autocomplete the rest. In this case, the dot and slash tell BASH to look in the current directory for any executable programs or scripts (programs with `X` as part of their permissions) and use them as possible autocomplete options.

BASH is clever enough to spot whether the command you're using is likely to require a file, directory, or executable, and it will autocomplete with only relevant file or directory names.

Viewing Available Options

The autocomplete function has a neat side effect. As I mentioned earlier, if BASH cannot find a match, pressing Tab again causes BASH to show all the available options. For example, typing `ba` at the shell, and then pressing Tab twice will cause BASH to show all the possible commands starting with the letters `ba`. On my test PC, this produces the following list of commands:

<code>badblocks</code>	<code>base-config</code>	<code>bash</code>	<code>batch</code>
<code>banner</code>	<code>basename</code>	<code>bashbug</code>	

This can be a nice way of exploring what commands are available on your system. You can then use each command with the `--help` command option to find out what it does, or browse the command's man page.

When you apply this trick to directory and filename autocompletion, it's even more useful. For example, typing `cd` in a directory, and then pressing the Tab key twice will cause BASH to show the available directories, providing a handy way of retrieving a brief directory listing. Alternatively, if you've forgotten how a directory name is spelled, you can use this technique to find out prior to switching into it.

Figure 17-1 shows a few examples of using this technique with BASH.

```

keir@ubuntu: /
File Edit View Terminal Tabs Help
keir@ubuntu:/$ ba
badblocks  base-config  bash          batch
banner     basename     bashbug
keir@ubuntu:/$ cd /
bin/       dev/         initrd.img    mnt/         sbin/        usr/
boot/     etc/         lib/          opt/         srv/         var/
cdrom/    home/        lost+found/   proc/        sys/         .viminfo
debootstrap/  initrd/    media/        root/        tmp/         vmlinuz
keir@ubuntu:/$ cd /home/
frank/ keir/
keir@ubuntu:/$ cd /home/keir/
.bash_history  .gconf/          .metacity/      .thumbnails/
.bash_profile  .gconfd/         .mozilla/       .Trash/
.bashrc        .gksu.lock       myfile          .update-notifier/
Desktop/       .gnome2/         .nautilus/     .viminfo
.dmrc         .gnome2_private/.openoffice.org2/.Xauthority
.esd_auth     .gstreamer-0.8/  README         .xsession-errors
.evolution/   .gtkrc-1.2-gnome2 README.zip
.fonts.cache-1 .ICEauthority    .recently-used
keir@ubuntu:/$ cd /home/keir/

```

Figure 17-1. Autocompletion makes using BASH much easier.

Using Keyboard Shortcuts

Your other good friends when using BASH are the Ctrl and Alt keys. These keys provide shortcuts to vital command-line shell functions. They also let you work more efficiently when typing by providing what most programs call keyboard shortcuts.

Shortcuts for Working in BASH

Table 17-1 lists the most common keyboard shortcuts in BASH (there are many more; see BASH's man page for details). If you've explored the Emacs text editor, you might find these shortcuts familiar. Such keyboard shortcuts are largely the same across many of the software packages that originate from the GNU Project. Often, you'll find an option within many Ubuntu software packages that lets you use Emacs-style navigation, in which case, these keyboard shortcuts will most likely work equally well.

Table 17-1. Keyboard Shortcuts in BASH

Shortcut	Description
Navigation	
Left/right cursor key	Move left/right in text
Ctrl+A	Move to beginning of line
Ctrl+E	Move to end of line
Ctrl+right arrow	Move forward one word
Ctrl+left arrow	Move left one word

Table 17-1. *Keyboard Shortcuts in BASH (Continued)*

Shortcut	Description
Editing	
Ctrl+U	Delete everything behind cursor to start of line
Ctrl+K	Delete from cursor to end of line
Ctrl+W	Delete from cursor to beginning of word
Alt+D	Delete from cursor to end of word
Ctrl+T	Transpose characters on left and right of cursor
Alt+T	Transpose words on left and right of cursor
Miscellaneous	
Ctrl+L	Clear screen (everything above current line)
Ctrl+U	Undo everything since last command ¹
Alt+R	Undo changes made to the line ²
Ctrl+Y	Undo deletion of word or line caused by using Ctrl+K, Ctrl+W, and so on ³
Alt+L	Lowercase current word (from the cursor to end of word)

¹ In most cases, this has the effect of clearing the line.

² This is different from Ctrl+U, because it will leave intact any command already on the line, such as one pulled from your command history.

³ This allows primitive cutting and pasting. Delete the text and then immediately undo, after which the text will remain in the buffer and can be pasted with Ctrl+Y.

Shortcuts for System Control

In terms of the control over your system offered by keyboard commands, as mentioned in Chapter 16, pressing Ctrl+Z has the effect of stopping the current program. It suspends the program until you switch back into it or tell it to resume in another way, or manually kill it.

In the same style, pressing Ctrl+C while a program is running will quit it. This sends the program's process a termination signal, a little like killing it using the top program. Ctrl+C can prove handy if you start a program running by accident and quickly want to end it, or if a command takes longer than you expected to work and you cannot wait for it to complete. It's also a handy way of attempting to end crashed programs. Some complicated programs don't take too kindly to being quit in this way, particularly those that need to save data before they terminate. However, most should be okay.

Ctrl+D is another handy keyboard shortcut. This sends the program an end-of-file (EOF) message. In effect, this tells the program that you've finished your input. This can have a variety of effects, depending on the program you're running. For example, pressing Ctrl+D on its own at the shell prompt when no program is running will cause you to log out (if you're using a GUI terminal emulator like GNOME Terminal, the program will quit). This happens because pressing Ctrl+D informs the BASH shell program that you've finished your input. BASH then interprets this as the cue that it should log you out. After all, what else can it do if told there will be no more input?

While it might not seem very useful for day-to-day work, Ctrl+D is vital for programs that expect you to enter data at the command line. You might run into these as you explore BASH. If ever you read in a man page that a program requires an EOF message during input, you'll know what to press.

Using the Command History

The original hackers who invented the tools used under Unix hated waiting around for things to happen. After all, being a hacker is all about finding the most efficient way of doing any particular task.

Because of this, the BASH shell includes many features designed to optimize the user experience. The most important of these is the *command history*. BASH remembers every command you enter (even the ones that don't work!) and stores them as a list on your hard disk.

During any BASH session, you can cycle through this history using the up and down arrow keys. Pressing the up arrow key takes you back into the command history, and pressing the down arrow key takes you forward.

The potential of the command history is enormous. For example, rather than retype that long command that runs a program with command options, you can simply use the cursor keys to locate it in the history and press Enter.

■ **Tip** Typing `!-3` will cause BASH to move three paces back in the history file and run that command. In other words, it will run what you entered three commands ago.

On my Ubuntu test system, BASH remembers 1000 commands. You can view all of the remembered commands by typing `history` at the command prompt. The history list will scroll off the screen because it's so large, but you can use the scroll bars of the GNOME Terminal window to read it.

Each command in the history list is assigned a number. You can run any of the history commands by preceding their number with an exclamation mark (!), referred to as a *bang*, or sometimes a *shriek*. For example, you might type `!923`. On my test system, command number 923 in the BASH history is `cd ..`, so this has the effect of switching me into the parent directory.

Command numbering remains in place until you log out (close the GNOME Terminal window or end a virtual console session). After this, the numbering is reordered. There will still be 1000 commands, but the last command you entered before logging out will be at the end of the list, and the numbering will work back 1000 places until the first command in the history list.

■ **Tip** One neat trick is to type two bangs: `!!`. This tells BASH to repeat the last command you entered.

Rather than specifying a command number, you can type something like `!cd`. This will cause BASH to look in the history file, find the last instance of a command line that started with `cd`, and then run it.

Pressing Ctrl+R lets you search the command history from the command prompt. This particular tool can be tricky to get used to, however. As soon as you start typing, BASH will autocomplete the command based on matches found in the history file, starting with the last command in the history. What you type appears before the colon, while the autocompletion appears afterwards.

Because BASH autocompletes as you type, things can get a little confusing when you're working with the command history, particularly if it initially gets the match wrong. For example, typing `cd` will show the last instance of the use of `cd`, as in the example in Figure 17-2. This might not be what you're looking for, so you must keep typing the command you do want until it autocompletes correctly.



Figure 17-2. BASH history completion is very useful but can also be confusing.

Piping and Directing Output

It's not uncommon for a directory listing or output from another command to scroll off the screen. When using a GUI program like GNOME Terminal, you can use the scroll bars to view the output, but what if you are working at the bare command-line prompt?

By pressing Shift+Page Up and Shift+Page Down, you can “scroll” the window up to take a look at some of the old output, but very little is cached in this way, and you won't see more than a few screens. A far better solution is to pipe the output of the directory listing into a text viewer. Another useful technique is to redirect output to a file.

Piping the Output of Commands

Piping was one of the original innovations provided by Unix. It simply means that you can pass the output of one command to another, which is to say the output of one command can be used as input for another.

This is possible because shell commands work like machines. They usually take input from the keyboard (referred to technically as *standard input*) and, when they've done their job, usually show their output on the screen (known as *standard output*).

The commands don't need to take input from the keyboard, and they don't need to output to the screen. Piping is the process of diverting the output before it reaches the screen and passing it to another command for further processing.

Let's assume that you have a directory that is packed full of files. You want to do a long directory listing (`ls -l`) to see what permissions various files have. But doing this produces reams of output that fly off the screen. Typing something like the following provides a solution:

```
ls -l | less
```

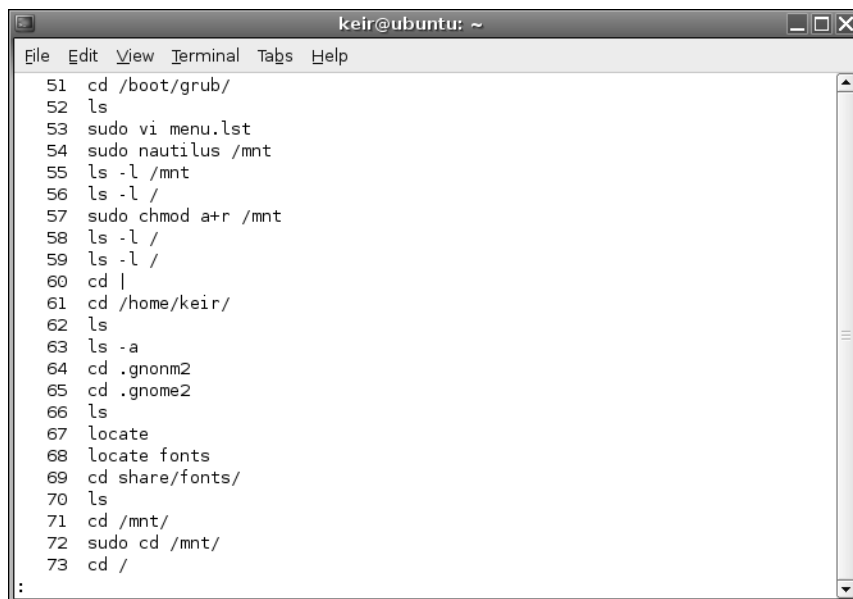
The `|` symbol between the two commands is the pipe. It can be found on most US keyboards next to the square bracket keys (above the Enter key; you'll need to hold down the Shift key to get it).

What happens in the example is that `ls -l` is run by the shell, but rather than sending the output to the screen, the pipe symbol (`|`) tells BASH to send it to the command that follows—to `less`. In other words, the listing is displayed within `less`, where you can read it at your leisure. You can use Page Up and Page Down or the arrow keys to scroll through it. Once you quit `less`, the listing evaporates into thin air; the piped output is never actually stored as a file.

In the previous section, you saw how you can use the `history` command to view the command history. At around 1,000 entries, its output scrolls off the screen in seconds. However, you can pipe it to `less`, like so:

```
history | less
```

Figure 17-3 shows the result on my test PC.



```
keir@ubuntu: ~
File Edit View Terminal Tabs Help
51 cd /boot/grub/
52 ls
53 sudo vi menu.lst
54 sudo nautilus /mnt
55 ls -l /mnt
56 ls -l /
57 sudo chmod a+r /mnt
58 ls -l /
59 ls -l /
60 cd |
61 cd /home/keir/
62 ls
63 ls -a
64 cd .gnom2
65 cd .gnome2
66 ls
67 locate
68 locate fonts
69 cd share/fonts/
70 ls
71 cd /mnt/
72 sudo cd /mnt/
73 cd /
:
```

Figure 17-3. Piping the output of the `history` command into the `less` command lets you read the output fully.

You can pipe the output of any command. One of the most common uses is when searching for a particular string in the output of a command. For example, let's say you know that, within a crowded directory, there's a file with a picture of some flowers. You know that the word *flower* is in the filename, but can't recall any other details. One solution is to perform a directory listing,

and then pipe the results to `grep`, which is able to search through text for a user-defined string (see Chapter 15):

```
ls -l | grep -i 'flower'
```

In this example, the shell runs the `ls -l` command, and then passes the output to `grep`. The `grep` command then searches the output for the word *flower* (the `-i` option tells it to ignore uppercase and lowercase). If `grep` finds any results, it will show them on your screen.

The key point to remember is that `grep` is used here as it normally is at the command prompt. The only difference is that it's being passed input from a previous command, rather than being used on its own.

You can pipe more than once on a command line. Suppose you know that the filename of the picture you want involves the words *flower* and *daffodil*, yet you're unsure of where they might fall in the filename. In this case, you could type the following:

```
ls -l | grep -i flower | grep -i daffodil
```

This will pass the result of the directory listing to the first `grep`, which will search the output for the word *flower*. The second pipe causes the output from `grep` to be passed to the second `grep` command, where it's then searched for the word *daffodil*. Any results are then displayed on your screen.

Redirecting Output

Redirecting is like piping, except that the output is passed to a file rather than to another command. Redirecting can also work the other way: the contents of a file can be passed to a command.

If you wanted to create a file that contained a directory listing, you could type this:

```
ls -l > directorylisting.txt
```

The angle bracket (>) between the commands tells BASH to direct the output of the `ls -l` command into a file called `directorylisting.txt`. If a file with this name exists, it's overwritten with new data. If it doesn't exist, it's created from scratch.

You can add data to an already existing file using two angle brackets:

```
ls -l >> directorylisting.txt
```

This will append the result of the directory listing to the end of the file `directorylisting.txt`, although, once again, if the file doesn't exist, it will be created from scratch.

Redirecting output can get very sophisticated and useful. Take a look at the following:

```
cat myfile1.txt myfile2.txt > myfile3.txt
```

As you learned in Chapter 15, the `cat` command joins two or more files together. If the command were used on its own without the redirection, it would cause BASH to print `myfile1.txt` on the screen, immediately followed by `myfile2.txt`. As far as BASH is concerned, it has joined `myfile1.txt` to `myfile2.txt`, and then sent them to standard output (the screen). By specifying a redirection, you have BASH send the output to a third file. Using `cat` with redirection is a handy way of combining two files.

It's also possible to direct the contents of a file back into a command. Take a look at the following:

```
sort < textfile.txt > sortedtext.txt
```

The `sort` command simply sorts words into alphanumeric order (it actually sorts them according to the ASCII table of characters, which places symbols and numbers before alphabetic characters). Directly after the `sort` command is a left angle bracket, which directs the contents of the file specified immediately after the bracket into the `sort` command. This is followed by a right angle bracket, which directs the output of the command into another file.

Tip To see a table of the ASCII characters, type `man 7 ascii` at the command-line prompt.

There aren't many instances in day-to-day usage where you'll want to use the left angle bracket. It's mostly used with the text-based `mail` program (which lets you send e-mail from the shell), and in shell scripting, in which a lot of commands are combined together to form a simple program.

REDIRECTING STANDARD ERROR OUTPUT

Standard input and standard output are what BASH calls your keyboard and screen. These are the default input and output methods that programs use unless you specify something else, such as redirecting or piping output and input.

When a program goes wrong, its error message doesn't usually form part of standard output. Instead, it is output via *standard error*. Like standard output, this usually appears on the screen.

Sometimes, it's very beneficial to capture an error message in a text file. This can be done by redirecting the standard error output. The technique is very similar to redirecting standard output:

```
cdrecord --scanbus 2> errormessage.txt
```

The `cdrecord` command is used to burn CDs, and with the `--scanbus` command option, you tell it to search for CD-R/RW drives on the system, something which frequently results in an error message if your system is not properly configured.

After the initial command, you see the redirection. To redirect standard error, all you need to do is type `2>`, rather than simply `>`. This effectively tells BASH to use the second type of output: standard error.

You can direct both standard output and standard error to the same file. This is done in the following way:

```
cdrecord --scanbus > error.txt 2>&1
```

This is a little more complicated. The standard output from `cdrecord --scanbus` is sent to the file `error.txt`. The second redirect tells BASH to include standard error in the standard output. In other words, it's not a case of standard output being written to a file, and then standard error being added to it. Instead, the standard error is added to standard output by BASH, and then this is written to a file.

Summary

In this chapter, we've looked at some tricks and tips to help you use the BASH shell more effectively. You've seen how BASH can help by autocompleting commands, filenames, and directories. You also learned about keyboard shortcuts that can be used to speed up operations within the shell.

This chapter also covered the command history function and how it can be used to reuse old commands, saving valuable typing time. Finally, we looked at two key functions provided by BASH: redirection and piping. This involved the explanation of standard input, output, and error.

In Part 5 of the book, starting with the next chapter, we move on to discuss the multimedia functionality within Ubuntu.

PART 5



Multimedia



Digital Music

Today's PC is a multimedia powerhouse, and it's hard to come across a home computer that doesn't have at least a set of speakers attached. Some people take this to extremes and have surround-sound speakers on their computers, as well as large monitors for crystal-clear video playback.

The people behind Ubuntu aren't blind to this and include not only an audio player, but also a video player with the distribution. In this chapter, you'll learn how to listen to MP3s, CDs, and Internet radio on your Ubuntu system. In the next chapter, you'll learn how to manage video playback.

Understanding the Legalities of Playback

As you might have read in the press, multimedia playback on computer devices, and Linux in particular, is hindered by a number of issues, including the following:

Software patents: Audio and video playback technologies such as MP3 and MPEG are patented in countries that allow software to be patented, such as the United States. A *patent* protects the implementation of an idea, as opposed to *copyright*, which protects the actual software. Patents are designed to restrict use of the technology unless permission is granted, usually via a payment to the license holder. Because Linux is based on the sharing of computing technology and knowledge, organizations like Ubuntu are fundamentally and philosophically opposed to any kind of software patenting. While this doesn't make playback of popular music and video files impossible under Ubuntu, it does mean that extra software must be downloaded and installed. Furthermore, depending on where you live in the world, there may be legal issues surrounding using such patented software.

Digital Rights Management (DRM): Much more devastating than patenting is DRM, a technology tied into audio or video playback software. It's designed to control how, where, when, and on what device you can play certain media. For example, Apple's iTunes DRM scheme means you can play back MP3s bought from iTunes only on their iPod range of devices (including the Motorola Rokr phone) or using the iTunes software. DVD movie players include a form of DRM called Content Scrambling System (CSS), which prevents users from playing DVDs on computers unless special software is purchased. Perhaps it goes without saying that the Linux community, including the Ubuntu project, is also fundamentally opposed to DRM. Because of this, practically no DRM software has been officially ported to Linux, so you can't, for example, play music purchased via the iTunes or Napster online stores.

Linux and other open-source programmers are very resourceful and are often able to reverse-engineer technology formats in order to get around DRM or patent issues. But the laws in many countries—with the United States as a particularly strident example—prohibit reverse engineering in this way. In addition, the laws in some countries seek to prohibit use of software resulting from this process.

Note You may be wondering why music and movie corporations are so intent on enforcing DRM and patenting if these schemes give their customers such a hard time. To learn more, and to find out what you can do to help halt the progress of such technology, visit the Electronic Frontier Foundation’s website: www.eff.org.

Programmers have also come up with Free Software alternatives to proprietary formats. Examples include the Ogg media format, which is every bit as good as MP3, but is unencumbered by patent issues. We’ll look at using Ogg later in this chapter, in the “Choosing a Format” section. But at the moment, there’s no ideal open-source video format, or at least not one that’s in widespread use.

Therefore, you’ll probably want to add support for MP3 and other popular music file formats. Throughout this and the next chapter, we’ll examine installing media playback software, some of which may have issues surrounding patenting. In one case, the software is designed to bypass the DRM scheme that protects DVD movie discs.

Note The United States and Japan both have laws allowing software to be patented. Most other countries, including those within the European Union, do not currently allow software patents.

Playing Music Files

Music playback under Ubuntu is handled by the Rhythmbox player. Like many modern music players, Rhythmbox can also manage your music collection, arranging it into a library so you can locate songs easily.

But before using Rhythmbox, you need to add support for MP3 and other popular music file formats.

Installing Codecs

The piece of software that handles the decoding (and also encoding) of digital music files is called a *codec*. The word is a shortened version of *coder-decoder*. For any digital multimedia file type you want to play back on your computer, you’ll need an appropriate codec. This includes both audio and video files. In addition, if you wish to create your own multimedia files, you might need to download an additional codec that allows the *encoding* of files.

The necessary audio and video codec software for playback can be found in Ubuntu's online software repositories. (See Chapter 8 for details on setting up the Synaptic Package Manager to access these repositories.) To obtain them, open the Synaptic Package Manager (Settings ► Administration), click the Search button, and then type `gststreamer`. In the list of results, look for the following packages:

- `gststreamerx.x-plugins`
- `gststreamerx.x-plugins-multiverse`
- `gststreamerx.x-ffmpeg`

The `x.x` represents the package version (0.8 was the latest at the time of writing). Look for the newest versions.

Click the check box alongside each and select Mark for Installation. You'll be told that each package needs additional software, as shown in Figure 18-1. Don't worry—this is why I recommend these particular packages! The set of packages contain support for just about every form of audio playback possible on a PC (with the exception of DRM technology like the iTunes `.m4p` files).

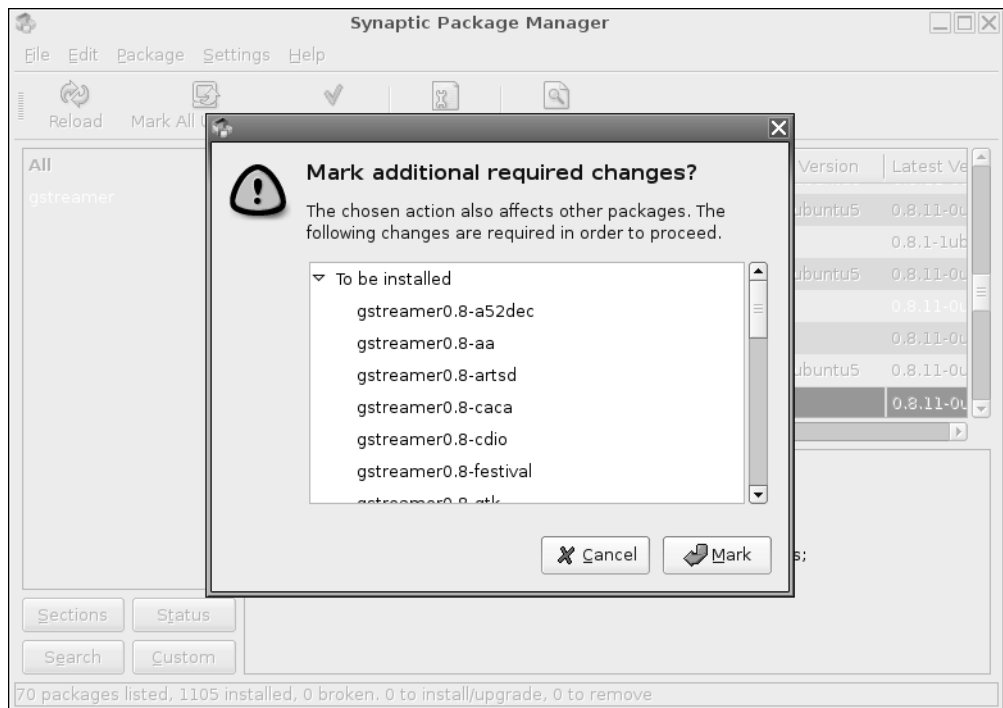


Figure 18-1. Installing the `gststreamer` packages will install all the audio codecs you need.

Click the Apply button, and the Synaptic Package Manager will download and install the software. Once this task has completed, quit the program.

Using Rhythmbox

You'll find Rhythmbox on the Applications ► Sound & Video menu. The first time you run the program, it will ask you where your music files are stored. Simply click the Browse button, and then double-click the folder that holds those files. Then click OK.

When the program starts, your music files are listed at the bottom of the program window. At the top left of the program window, you'll find a listing of the artists behind the MP3s in your collection, and on the right, you'll see the album that the music track is taken from (provided that information is included in the music file itself, such as the MP3 ID3 tags). Figure 18-2 shows an example of a Rhythmbox window.

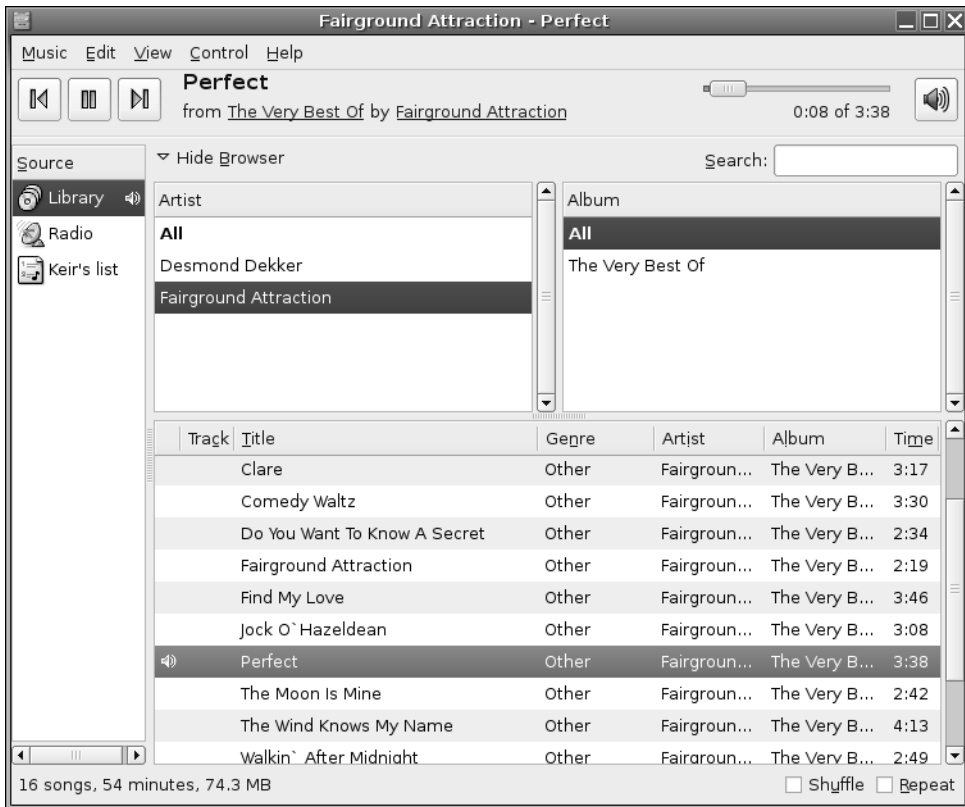


Figure 18-2. *Rhythmbox will organize your music tracks by artist or album.*

Playing a track is simply a matter of double-clicking it in the list. Once the track is finished, Rhythmbox will play the next track in the MP3 file list.

To create a playlist, right-click under the Source heading on the left side of the program window and select New Playlist. Click the new entry to type its name. To add tracks to the playlist, click Library, and then drag-and-drop files onto your new playlist entry. To start the playlist, click the first track in its list.

Note I noticed that whenever I double-clicked an MP3 in a Nautilus window, Ubuntu would try to start the Totem Movie Player rather than Rhythmbox. This was easy to fix: right-click any MP3 file, select Properties, and click the Open With tab. Then click the Add button, select Rhythmbox, and ensure that the radio button alongside its entry in the list is selected.

Listening to Audio CDs

Playing back audio CDs is simple. Just insert the CD, and then click Applications ► Sound & Video ► CD Player. This will start Ubuntu's simple but effective CD playback application, as shown in Figure 18-3. If you're online, the application will immediately attempt to look up the CD artist and track listing information from an online repository. This means that it should show the names of the tracks being played.

Note On one occasion during my testing, the online lookup resulted in several entries for a CD I was playing. All I had to do was choose the one that I was happiest with (not hard considering the other two were in German!). From then on, whenever I played that CD, the track listing I had chosen was used.



Figure 18-3. *Ubuntu's CD Player application is simple but effective and can look up artist and track information online.*

The controls work in the same way as those on any other audio player. Clicking the Play button plays the current track, which you can select from the drop-down list above the player controls. You can skip to the next track by clicking the Next Track button and cue backwards and forwards using the relevant buttons. The time display in the top left of the program window shows the track's progress.

Tip As with all GNOME applications, hover the mouse cursor over each button to display a tooltip that describes what it does.

If you find the track listing information is incorrect, you can correct it by clicking the Track Editor button on the right side of the program window. In the window that appears, double-click the track entry with the incorrect information and edit it as you see fit.

Tuning in to Online Radio Stations

You can listen to online radio stations (known as streaming audio) by downloading the Streamtuner software. This isn't an actual audio player, but is instead a kind of digital listings program that lets you search for stations to match just about any musical taste. It relies on the XMMS media player program to actually allow you to listen to the streams.

To obtain the software, open Synaptic Package Manager (Settings ► Administration), search for `streamtuner`, and mark it for installation. Then search for `xmms` and mark it for installation, too. Click Apply to install all the packages.

Configuring XMMS

Before using XMMS, you'll need to quickly reconfigure its sound output plug-in so that its volume control works properly. Here are the steps:

1. Start XMMS by clicking its entry on the Applications ► Sound & Video menu.
2. Right-click its title bar and select Options ► Preferences.
3. In the dialog box that appears, look for the Output Plugin heading under the Audio I/O Plugins tab, and select eSound Output Plugin.
4. Click OK, and then quit XMMS.

On my installation of Ubuntu, I also had to create a desktop shortcut for Streamtuner because, for some reason, it wasn't added to the Sound & Video submenu. If you need to do this, right-click a blank spot on the desktop and select Create Launcher. In the Name box, enter **Streamtuner**, and in the Command box, enter **streamtuner**. Click to add an icon, and then enter the following in the Path field of the icon selection dialog box:

```
/usr/share/streamtuner/python/icons/basic.ch.png.
```

Using Streamtuner

Using Streamtuner is easy. When it launches, click one of the tabs to browse through the stations. In addition, some tabs have subcategories of stations on the left. The Shoutcast tab probably contains the most comprehensive selection, but some people prefer Live365, shown in Figure 18-4.

Double-click a stream to start it playing. Most streams will play in XMMS, which will start automatically, but some are RealAudio streams, and you'll need to have RealPlayer installed to play them. See Chapter 19 for instructions on installing RealPlayer.

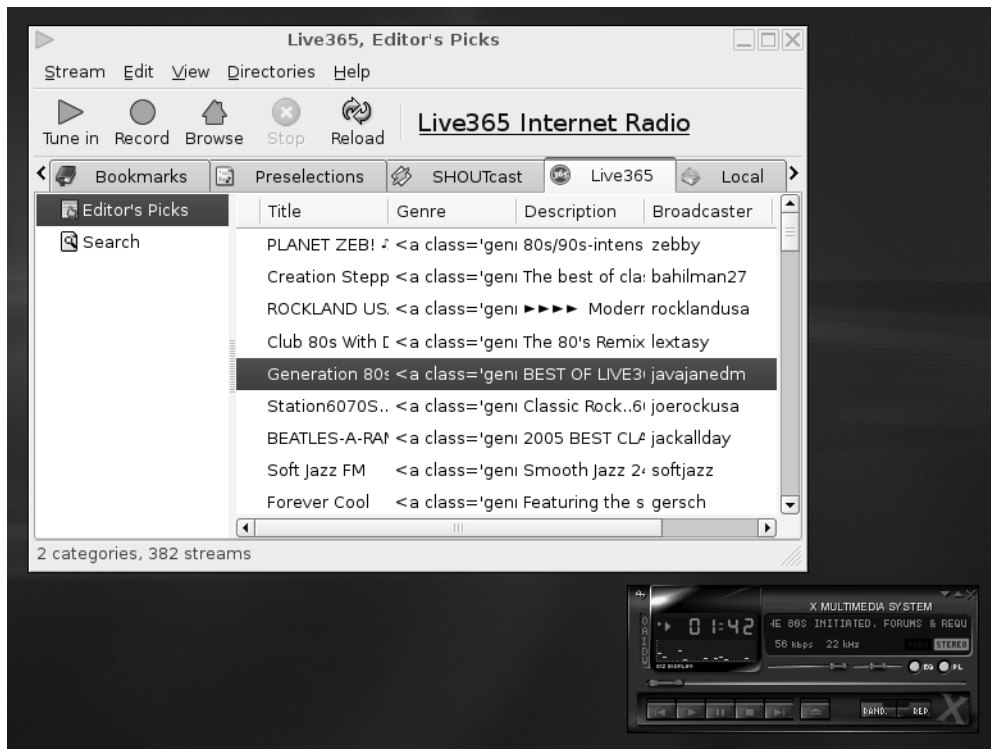


Figure 18-4. With *Streamtuner* and *XMMS*, you can tune into a massive range of Internet radio stations.

Ripping Music from CDs

Converting audio tracks on a CD into digital music files you can store on your hard disk for personal use is informally known as *ripping*. It's handled under Ubuntu using the Sound Juicer application, which can be found on the Applications ► Sound & Video menu.

Note Because of the way audio CDs work, you can't simply insert the disc and then drag-and-drop the tracks onto your hard disk. They must be converted first.

Before you start to rip CDs, however, you'll need to decide the format in which you wish to store the audio files.

Choosing a Format

You have three basic choices for audio file formats: *Ogg Vorbis*, *FLAC*, and *MP3*. Let's look at what each has to offer.

Ogg Vorbis: This is the Free Software alternative to MP3. Unless you have a trained ear, you won't be able to tell the difference between an Ogg and MP3 file (if you *do* have a trained ear, then you'll find Ogg better!). The two technologies generate files of around the same size, an average of 4MB to 5MB per song. The downside of Ogg is that not many portable audio players support it (although this situation is slowly changing), and other operating systems like Windows won't be able to play back Ogg files unless some additional software is installed (see www.vorbis.com/setup/).

FLAC: This stands for Free Lossless Audio Codec, and it's the choice of the audiophile. Ogg and MP3 are lossy formats, which means that some of the audio data is lost in order to significantly shrink the file. FLAC doesn't lose any audio data but still manages to compress files to a certain degree (although they're still much larger than an equivalent MP3 or Ogg file). FLAC scores points because it's Free Software, like Ogg, but you'll face the same lack of support in portable audio players and other operating systems (unless additional software is installed; see <http://flac.sourceforge.net/download.html>).

MP3: This is by far the most ubiquitous music file format and practically everyone who owns a computer has at least a handful of MP3 tracks. This means software support for MP3 playback is strong and, of course, portable audio players are built around the MP3 format. The only problem for you, as a Linux user, is the issue of surrounding patents, as explained at the beginning of this chapter. Using the MP3 format is to go against a lot of what the Free Software movement stands for. But in the end, the choice is up to you.

Adding MP3 Support to Sound Juicer

Support for Ogg and FLAC is built into Ubuntu, but if you wish to encode CD tracks as MP3s, you'll need to enter some configuration details into Sound Juicer.

Note The software that's required to encode MP3s is installed with the `gststreamer` packages. If you didn't already download them, follow the instructions earlier in the chapter, in the "Installing the Codecs" section.

Follow these steps to configure MP3 support:

1. Select Applications ► Sound & Video ► Sound Juicer to start the program.
2. Click Edit ► Preferences. In the dialog box that appears, click the Edit Profiles button.
3. Click New, and then type **MP3**. Then click Create.
4. Select the new MP3 entry in the list and click Edit.
5. In the Profile Name box, type **MP3**.
6. In the Profile Description box, type **MP3**.
7. In the GStreamer Pipeline box, type the following:

```
audio/x-raw-int,rate=44100,channels=2 ! lame name=enc
```

8. In the File Extension box, type **mp3**.
9. Put a check in the Active? box.
10. Click OK.
11. Restart Sound Juicer.
12. Click Edit ► Preferences again, and select your new MP3 entry from the Output drop-down list.

Ripping Tracks

When you're ready to rip some music, insert the audio CD, and then start Sound Juicer. If the CD isn't read immediately, click Disc ► Re-read. As with CD Player, Sound Juicer will automatically look up the artist and track information of the CD.

Click Edit ► Preferences. In the Format part of the dialog box, choose the type of audio files you want to create. In addition, you can select where you would like the files to be saved to by clicking the drop-down list under the Music Folder heading.

Any track in the listing with a check in its box will be ripped. When you insert a CD, Sound Juicer assumes that all the tracks are to be ripped, as shown in Figure 18-5. If this isn't the case, remove the checks from the tracks you don't want to rip. By selecting a track and clicking the Play button, you'll be able to audition the track. This can be helpful if you're deciding on exactly which tracks to rip. Finally, check that the Title and Artist information is correct.

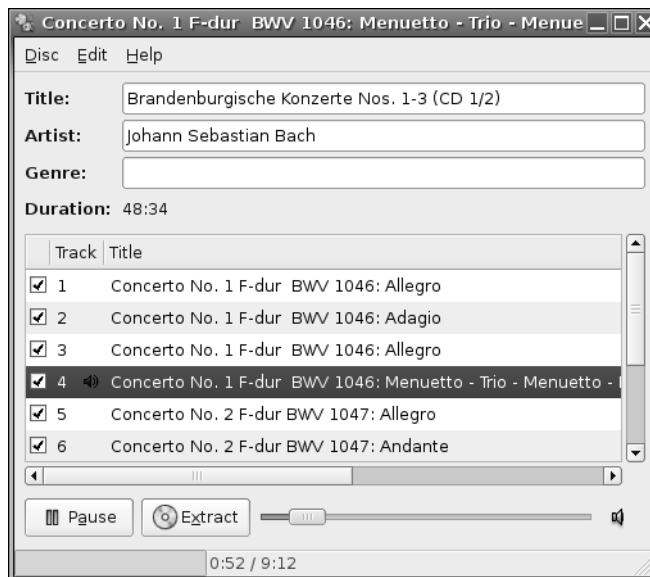


Figure 18-5. Audio tracks can be ripped from CDs using the Sound Juicer program.

To begin the process, click the Extract button. It will take up to a minute or so to rip each track, so ripping an entire CD may take some time. However, it's safe to leave Sound Juicer working in the background.

MAKING MUSIC AND RECORDING AUDIO

Most PCs come with sound cards that are capable of making music. You can use many open-source programs, designed for both amateurs and professionals alike, to create music or record and edit audio.

In terms of musical sequencers, Muse (<http://muse.serverkommune.de>), Rosegarden (www.rosegardenmusic.com), and Jazz++ (www.jazzware.com/zope) are well worth investigating. Like all modern MIDI sequencers, all three programs let you record audio tracks, effectively turning your PC into a recording studio.

It's also possible to run virtual synthesizers on your PC, which effectively turn even the most basic sound card into a powerful musical instrument. Examples include Bristol (www.slabexchange.org) and FluidSynth (www.fluidsynth.org).

If you're interested in only audio recording and processing, Sweep (www.metadecks.org/software/sweep/) and Audacity (<http://audacity.sourceforge.net>) are worth a look. In addition to audio recording and playback, both feature graphical waveform editing and powerful filters.

Most of the packages mentioned here are available from the Ubuntu software repositories, and you can download them with the Synaptic Package Manager.

Creating Your Own CDs

You can create audio CDs by using the Serpentine Audio-CD Creator program, found on the Applications ► Sound & Video menu.

Start by inserting a blank CD. Almost immediately, Ubuntu will ask what you want to do with the disc via a dialog box. If you choose Burn an Audio CD, Serpentine will open automatically.

The program is very simple to use. Click the Add button on the toolbar, and then browse to your store of files (these can be MP3, Ogg, or any other file format supported by Ubuntu). Then select the tracks you wish to go onto the CD. Shift-click to select many tracks at once, as with Windows, or hold down Ctrl and click multiple individual tracks.

After clicking Open to add the tracks, you'll see the track listing build up in the Serpentine program window. In addition, a small graphic of CD will also fill up, like a pie chart, showing how much space is left on the CD. You can rearrange the track listing by clicking and dragging the tracks to new locations.

When you're happy with the track listing, click the Write to Disc button to start the write procedure. First, the tracks are converted to pure audio files, and then they're actually burned to disc, as shown in Figure 18-6. This can take some time. When Serpentine finishes with the burning, the CD will be ejected.

Tip Depending on the quality of the blank CD, you might not be able to write audio CDs at full speed. If this is the case, Serpentine will stop during the writing process with an error message. Click Edit ► Preferences on the menu and click the radio button next to Choose Writing Speed. Then enter a more conservative speed.

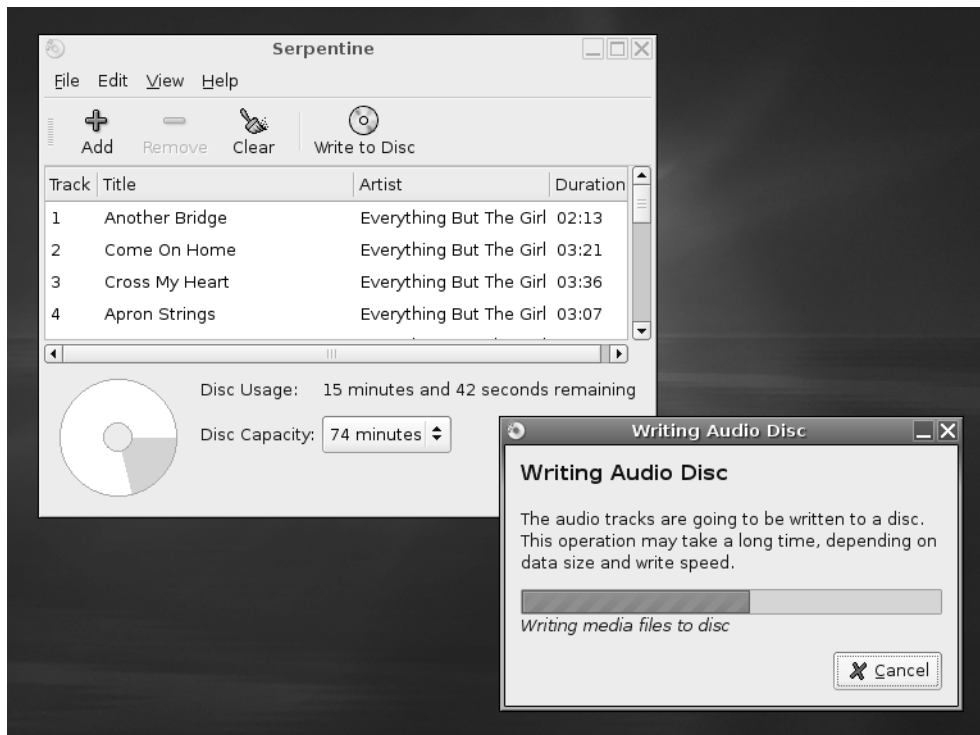


Figure 18-6. *Serpentine Audio-CD Creator makes it easy to create CDs from digital audio files.*

Summary

This chapter has covered the audio functions built into Ubuntu and how, by downloading a few extra system files, you can play back the majority of audio files in existence. We started by discussing the moral and legal dilemmas associated with multimedia playback on a computer. Then we moved on to look at how to install the necessary codec files on your computer, before discussing how you can listen to music files, CDs, and online radio stations.

We examined how you can convert CDs into music files, and then the inverse of this: how you can create CDs using audio files.

In Chapter 19, we'll look at how you can play back movies and online animations using Ubuntu.



Movies and Multimedia

Movie playback is becoming increasingly popular on computers. Modern PCs come equipped with DVD-ROM drives and, coupled to the right software, these can play DVD movie discs. In addition, many web sites feature streaming movie clips or offer them for download.

Ubuntu provides support for movie playback but, as with audio support, you'll need to install additional codecs in order to enjoy the broadest range of playback options. This chapter explains how to set up Ubuntu for watching videos, DVDs, and TV on your computer, as well as playing web site Flash animations.

Watching Videos

You use the Totem Movie Player application to play back video under Ubuntu. Like the other multimedia software provided with Ubuntu, it's basic but effective and does the job well. However, because of licensing issues, it doesn't support all video formats out of the box. In fact, it supports very few of those you might be used to using under Windows.

If you wish to play back the most common video file types, such as those listed in Table 19-1, you must install additional software. This software contains codecs that allow you to play Windows Media Player, RealPlayer, QuickTime, and DivX movie files under Ubuntu. In addition, you'll need to make some updates to Totem.

Caution The codecs are quite literally lifted straight from a Windows installation. Therefore, you should install them only if you own a license for Windows and have also installed the relevant media player applications under Windows and agreed to their license agreements.

After following the instructions in this section, you will be able to play video files on your hard disk and also stream video from web sites within Firefox.

Table 19-1. *Popular Movie File Formats*

Format	Typical File Extensions	Web Site	Notes
Windows Media Player 9	.wmv, .wma, .asx, .asf	www.microsoft.com/windows/windowsmedia	Default format for Windows Media Player and, therefore, for most Windows users. Although it's possible to play Windows Media Player files under Ubuntu, you won't be able to play DRM-restricted files (those that rely on the download and installation of a certificate).
Real Video	.rm, .ram	www.real.com	By downloading the Win32 codec pack, you can play back Real Video files in Totem. However, you can also download a Linux version of RealPlayer. Both approaches are described in this chapter.
QuickTime	.mov, .qt	www.quicktime.com	QuickTime is Apple's default media format and has gained ground on both Windows and Macintosh computers. As with Windows Media Player, you won't be able to play DRM-restricted files.
DivX	.avi, .divx	www.divx.com	The DivX format is one of the most popular formats for those in the Internet community who like to encode their own movies. It's renown for its ability to shrink movies to very small sizes.

Installing Codecs

For video playback, you need the codecs described in the previous chapter, as well as a few more. If you haven't followed the instructions in the previous chapter and downloaded the various packages mentioned there, do that first. Then follow these steps to install the additional codecs required:

1. Select **System** ► **Administration** ► **Synaptic Package Manager**.
2. Click **Search** and search for `totem-xine`. Click its check box in the results, mark it for installation, and then click **Apply**.
3. Select **Applications** ► **Accessories** ► **Terminal** and, in the GNOME Terminal window, enter the following in sequence, pressing **Enter** after each line (the **►** means to continue to the next line before pressing **Enter**):

```
wget ftp://ftp.nerim.net/debian-marillat/pool/main/w/w32codecs/w32codecs_►
20050412-0.0_i386.deb
sudo dpkg -i w32codecs_20050412-0.0_i386.deb
gst-register-0.8
```

4. Close the terminal window. Log out, and then log back in again.

Installing RealPlayer

If you wish to install RealPlayer, head over to www.real.com and follow the download links for the latest free Linux player, shown in action in Figure 19-1. Download it, open a GNOME Terminal window (Applications ► Accessories ► Terminal), and type the following (replacing *filename* with the name of the RealPlayer file):

```
chmod +x filename
sudo ./filename
```



Figure 19-1. A version of RealPlayer for Linux works just like its Windows brethren.

During installation, you'll be asked various questions. When asked where you would like to install RealPlayer, type `/usr/bin/RealPlayer` (don't agree to using the default path!). Answer yes to creating symbolic links, and agree to the default location for the links.

In my tests, I found it was necessary to restart the computer after installation to avoid unreliable playback in RealPlayer.

Playing Back Video

To play a movie file on your hard disk, simply double-click its icon. This will automatically start Totem, as shown in Figure 19-2.

Note By default, all video files will play in Totem, including RealMedia. To change this so that RealPlayer handles them, right-click any RealPlayer movie file, select **Open With**, and click **Add**. Locate RealPlayer in the list, click the **Add** button, and then make sure the radio button alongside RealPlayer is selected.



Figure 19-2. Properly configured, Totem can play just about every kind of movie file, such as QuickTime and Windows Media Player.

Using Totem is easy. At the bottom left of the screen are the transport controls that allow you to pause, play, or move forward and backward in the video file. Alternatively, you can right-click the video window and select the controls from there.

Above the controls is the Time bar. With certain types of video, you can drag-and-drop this to move through the video, but not all files support this function.

At the right of the program window is a playlist. You can queue several video files to be played in sequence by simply dragging-and-dropping movies from a Nautilus file browser window.

To play the video full screen, tap the F key. To return to the program window, press Escape. In full-screen mode, you can start and stop the video by pressing the spacebar.

You can hide the playlist by clicking the Sidebar button. This gives nearly all of Totem's program window to the playback window, but you might need to resize the window for the video to update to the new dimensions.

Tip If you find you have problems with video playback, such as Totem showing an error message about another application using the video output, try the following: click System ► Preferences ► Multimedia Systems Selector. Click the Video tab, and in the Default Sink/Output drop-down list, select Xwindows (No XV). Click Test. If this doesn't work, try the SDL – Simple DirectMedia Layer instead.

OPEN-SOURCE MOVIE FILE FORMATS

A number of promising open-source movie file formats are in development. Some are more mature than others, but few see widespread use at the moment. All promise much for the future. Many consider the following three formats as the chief contenders.

- XviD (www.xvid.org) is a reworking of the popular DivX MPEG-4-based file format. As such, it is able to encode movies to relatively small file sizes (a 90-minute movie can fit on a CD). Despite small file sizes, it can maintain good image and sound quality. In theory, it should also be possible to play XviD movies using any MPEG-4 codec, such as DivX or QuickTime. Unfortunately, XviD uses technology covered by patents in some parts of the world, so the project exists in a legally gray area. Additionally, it's only possible to download a Windows version of the codec, although if you follow the instructions at the beginning of this chapter, you will be able to download the ported Windows version of the codec so you can play XviD files under Ubuntu.
- Ogg Theora (www.theora.org) is being developed by the Xiph.org Foundation, the people behind the Ogg Vorbis audio codec project that's a favorite among Linux users. As such, it promises to be a completely open-source project. Although the technology is covered by patents, Xiph.org has promised never to enforce them, meaning that anyone in the world can use Theora without charge. At the time of writing, Theora is still in the alpha development phase, but it will almost certainly become the open-source video codec of choice in the future.
- The British Broadcasting Corporation (www.bbc.co.uk), the UK's largest public service broadcaster, is sponsoring development of the Dirac codec (see <http://dirac.sourceforge.net>). Dirac is less developed than both Theora and XviD at present, and it is aimed more at the broadcast/enthusiast market. For example, it is designed to support high-definition TV. However, it's certainly one to watch.

Watching DVDs

DVD movie discs are protected by a form of Digital Rights Management (DRM) called Content Scrambling System (CSS). This forces anyone who would like to create DVD playback software or hardware to pay a fee to the DVD Copy Control Association, an industry organization set up to protect DVD movie technology.

Nearly all Linux advocates are scornful of any kind of DRM system. It isn't possible to buy licensed DVD playback software for Linux but, even if it were, few would be willing to support what they see as prohibitive software technology.

Some open-source advocates reverse-engineered DVD and came up with the DeCSS software. This bypasses the CSS system and allows the playback of DVD movies under practically any operating system. Sadly, DeCSS is caught in a legal quagmire. The Motion Picture Association of America (MPAA) has attempted to stop its distribution within the United States but has failed. Some experts suggest that distributing DeCSS breaks copyright laws, but there has yet to be a case anywhere in the world that categorically proves this. Nor has there been a case proving or even suggesting that using DeCSS is in any way illegal.

Ubuntu doesn't come with DeCSS installed by default, but you can download and install the software by issuing a simple command. Here is the procedure:

1. Select System ► Administration ► Synaptic Package Manager.
2. Click Search and search for `libdvdread`. If the package isn't already installed, click its check box and mark it for installation. Then click Apply. Close Synaptic Package Manager.
3. Open a GNOME Terminal window (Applications ► Accessories ► Terminal). Type the following in the terminal window to download and install the DeCSS component:

```
sudo /usr/share/doc/libdvdread3/examples/install-css.sh
```

After you've installed DeCSS, just insert a DVD, and Totem will automatically start playing it. Alternatively, if Totem is already open, you can play the DVD by clicking its entry on the File menu.

Playing a DVD is not dissimilar to watching movie files on your hard disk. The only difference is that you can now navigate from chapter to chapter on the DVD by clicking the relevant entry under the Go menu, as shown in Figure 19-3. You can also return to the DVD's main menu this way. You can use your mouse to click entries in DVD menus.



Figure 19-3. Click the Go menu to access the DVDs menus and to skip forward or backward.

Playing Flash Animations

The Flash Player is a standard fixture on most modern browsers. It allows access to not only animations, but also to interactive web sites and games, as shown in Figure 19-4. Macromedia makes a player especially for Linux, and you can download it from the Ubuntu software repositories.




Figure 19-4. You can install Flash Player on Ubuntu to enjoy online animation, games, and interaction.

To install the Flash Player, open Synaptic Package Manager (System ► Administration) and search for `flashplayer-mozilla`. Then mark it for installation and click Apply. Once installation has finished, log out and then back in again. Configuration is automatic, and you should now be able to visit any Flash site with Firefox.

Tip If you have problems with sound playback when watching Flash animations, try the following: open a GNOME Terminal window (Applications ► Accessories ► Terminal), and then type `sudo ln -s /usr/lib/libesd.so.0 /usr/lib/libesd.so.1`. Then log out of Ubuntu and back in again.

Sadly, there isn't a Linux version of the Shockwave Director browser plug-in. If you really need to be able to have access to Shockwave sites under Linux, considering using CrossOver Office (www.codeweavers.com) to install the Windows version. But be aware that CrossOver Office is commercial product, and you'll need to pay for it.

Some sites use Java to present interaction, animation, and even movies. You can install the Java Runtime program by searching for `j2re` within the Synaptic Package Manager. In the list of results, click to mark for installation both the main `j2re` package and the `mozilla-plugin` package for `j2re`. When both are installed, log out and then back in again.

MOVIE EDITING

The field of Linux movie-editing software is still young, and only a handful of programs are available for the nonprofessional user. One of the best is Kino (www.kinodv.org), which is available in the Ubuntu software archives. Although far from being a professional-level program, Kino allows competent users to import and edit videos, apply effects, and then output in either MPEG-1 or MPEG-2 format.

If you're looking for something more powerful, but also more complicated, then Cinelerra is well worth a look (<http://heroinewarrior.com/cinelerra.php3>). To quote the web site, Cinelerra is "the same kind of compositing and editing suite that the big boys use," except it's made for Linux! Sadly, Cinelerra isn't in the Ubuntu repository and, at the time of writing, there wasn't a Ubuntu-compatible package at the web site. However, a sister project that is based on Cinelerra—Cinelerra-CV—contains a Debian package that should work fine on Ubuntu. See <http://cvs.cinelerra.org> for details.

MainActor (www.mainconcept.com) is the Linux version of a commercial Windows project. Although it's not free, most people agree that it's one of the most comprehensive video editors available for Linux at the moment, and possibly the easiest to use, too.

Incidentally, professional moviemakers use Linux all the time, particularly when it comes to adding special effects to movies. Movies like *Shrek 2*, *Stuart Little*, and the *Harry Potter* series all benefited from the CinePaint software running under Linux! For more details, see <http://cinpaint.movieeditor.com>.

Watching TV

If you have a TV card, you may be able to use it to watch TV under Ubuntu. Ubuntu doesn't come with a TV tuner application by default, but you can download the `tvtime` program from the software repositories using the Synaptic Package Manager.

Checking for Video Input

Ubuntu includes the Video for Linux project, an extension to the Linux kernel to allow many popular TV and video-capture cards to work. You can find out if yours is compatible by clicking **System** ► **Preferences** ► **Multimedia Systems Selector**. In the dialog box that appears, click the **Video** tab, and then click the **Test** button in the **Default Source** part of the window. If you see a video window without an error message, then your TV card is compatible. If you receive an error message, try a different input setting from the drop-down list and click **Test** again.

Installing Tvtime

To download and install `tvtime`, open Synaptic Package Manager (System ► Administration), click the Search button, and enter `tvtime` as a search term. In the list of results, click the entry for `tvtime` and mark it for installation. Then click Apply.

When the download has completed, you'll be asked a number of questions during the configuration process. First, you need to choose your TV picture format. Users in the United States should choose NTSC. Users in the United Kingdom, Australia, and certain parts of Europe should choose PAL. To find out which TV system your country uses, look up your country at www.videouniversity.com/standard.htm. You also need to choose your geographical area from the list so that `tvtime` can set the correct radio frequency range for your TV card.

Once the program is installed, you'll find it on the Applications ► Sound & Video menu. Figure 19-5 shows an example of `tvtime` in action. Using the program is straightforward, but if you need guidance, visit the program's web site at <http://tvtime.sourceforge.net>.



Figure 19-5. *If you have a TV tuner card, `tvtime` lets you watch TV on your PC!*

Summary

In this chapter, we looked how you can watch movies on your PC. You've seen how you can update Ubuntu to work with the most popular digital video technologies, such as Windows Media Player and QuickTime.

In addition, we looked at how you can view online multimedia such as Flash animations on your computer, and discussed how you can watch TV on your PC.

In the next chapter, we take a look at image editing under Ubuntu. You'll learn about one of the crown jewels of the Linux software scene: The GIMP.



Image Editing

The PC has become an increasingly useful tool in the field of photography. In fact, these days it's hard to imagine a professional photographer who doesn't use a computer in some way, either to download digital camera images or to scan in images taken using traditional film-based cameras.

Ubuntu includes a sophisticated and professional-level image-editing program called The GIMP. The title stands for GNU Image Manipulation Program. This chapter introduces this jewel in the crown of Linux software.

Getting Pictures onto Your PC

Before you can undertake any image editing, you need to transfer your images to your PC. Depending on the source of the pictures, there are a variety of ways of doing this. We have already looked at transferring images to your PC in Chapter 8, but let's briefly recap the procedure here.

Most modern cameras use memory cards to store the pictures. If you have such a model, when you plug the camera into your PC's USB port, you should find that Ubuntu instantly recognizes it. An icon should appear on the desktop, and double-clicking it should display the memory card in a Nautilus window. Technically speaking, the memory card has been *mounted* (see Chapter 14 for an explanation of mounting).

If your camera doesn't appear to be recognized by Ubuntu, you should consider buying a USB card reader. These devices are typically inexpensive and can read a wide variety of card types, making them a useful investment for the future. Some new PCs even come packaged with card readers. Most generic card readers should work fine under Linux, as will most new digital cameras.

If your camera isn't recognized, however, or if it's a few years old and uses the serial port to connect to your PC, you can try using the gThumb software (Applications ► Graphics ► gThumb Image Viewer), as explained in Chapter 8.

If you're working with print photos, negative film, or transparencies, you can use a scanner to scan them in using the XSane image scanning program, also covered in Chapter 8. This works in a virtually identical way to the TWAIN modules supplied with Windows scanners, in that you need to set the dots per inch (DPI) figures, as well as the color depth. Generally speaking, 300 DPI and 24-bit color should lead to a true-to-life representation of most photos (although because of their smaller size, transparencies or negative film will require higher resolutions, on the order of 1,200 or 2,400 DPI).

Introducing The GIMP

The GIMP is an extremely powerful image editor that offers the kind of functions usually associated with top-end software like Adobe Photoshop. Although it's not aimed at beginners, those new to image editing can get the most from of it, provided they put in a little work.

The program relies on a few unusual concepts within its interface, which can catch many people off guard. The first of these is that each of the windows within the program, such as floating dialog boxes or palettes, gets its own Panel entry. In other words, The GIMP's icon bar, image window, settings window, and so on have their own buttons on the Panel alongside your other programs, as if they were separate programs.

Note The GIMP's way of working is referred to as a Single Document Interface, or SDI. It's favored by a handful of programs that run under Linux and seems to be especially popular among programs that let you create things.

Because of the way that The Gimp runs, before you start up the program, it's a wise idea to switch to a different virtual desktop (virtual desktops are discussed in Chapter 7), which you can then dedicate entirely to The GIMP.

Click Applications ► Graphics ► GIMP Image Editor to run The Gimp. When the program starts for the first time, it will run through its setup routine. Usually, you can use the default answers to the various questions asked by the wizard.

After program has setup, you'll be greeted by what appears to be a complex assortment of program windows. Now you need to be aware of a second unusual aspect of the program: its reliance on right-clicking. Whereas right-clicking usually brings up a context menu offering a handful of options, within The GIMP, it's the principal way of accessing the program's functions. Right-clicking an image brings up a menu offering access to virtually everything you'll need while editing. Ubuntu includes the latest version of The GIMP, 2.2, and this features a menu bar in the main image-editing window. This is considered sacrilege by many traditional The GIMP users, although it's undoubtedly useful for beginners. However, the right-click menu remains the most efficient way of accessing The GIMP's tools.

The main toolbar window, shown in Figure 20-1, is on the left. This can be considered the heart of The GIMP because, when you close it, all the other program windows are closed, too. The menu bar on the toolbar window offers most of the options you're likely to use to start out with The GIMP. For example, File ► Open will open a browser dialog box in which you can select files to open in The GIMP. It's even possible to create new artwork from scratch by choosing File ► New, although you should be aware that The GIMP is primarily a photo editor. To create original artwork, a better choice is a program like OpenOffice.org Draw (on the Applications ► Office menu).

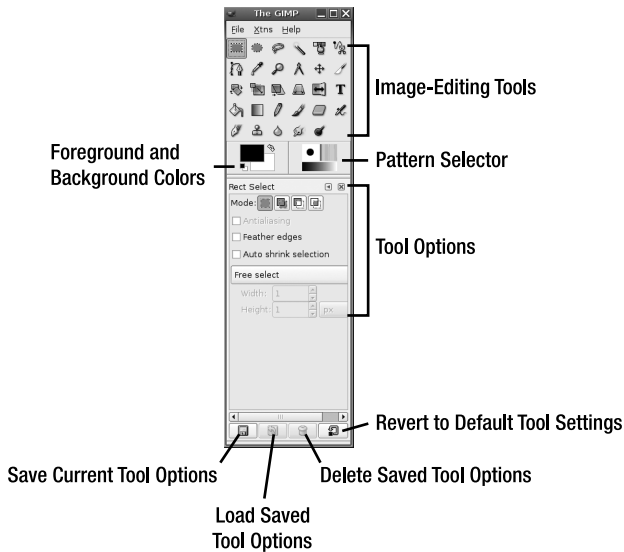


Figure 20-1. *The GIMP's main toolbar window*

Beneath the menu bar in the main toolbar window are the tools for working with images. Their functions are described in Table 20-1, which lists the tools in order from left to right, starting at the top left.

Table 20-1. *The GIMP Image-Editing Tools*

Tool	Description of Use
Rectangular selection tool	Click and drag to select a rectangular area within the image. This selected area can then be copied and pasted into a different part of the image or turned into a new layer.
Elliptical selection tool	Create an oval or circular selection area within the image, which you can then copy and paste.
Hand-drawn selection tool	Click and draw with the mouse to create a hand-drawn selection area. Your selection should end where it started. If not, The GIMP will draw a straight line between the start and end of the selection.
Contiguous regions selection tool	Known as the “magic wand” in other image editors, this tool creates a selection area based on the color of the pixels where you click. For example, clicking on a red car hood will select most, if not all of the hood, because it is mostly red.
Color region selection tool	This tool works like the contiguous region selection tool, but will create a selection across the entire image based on the color you select. In other words, selecting a black T-shirt will also select a black signpost elsewhere in the picture if the shades are similar.

Table 20-1. *The GIMP Image-Editing Tools (Continued)*

Tool	Description of Use
Shape selection tool	Another “magical” tool, the shape selector lets you create a selection by clicking on various points within an image, with the program joining the points together based on the color differences between the two points. This means that you can select the outline of a car by clicking a few points around the edge of the car and, provided the color of the car is different from the background, The GIMP will work out the color differences and select the car’s shape automatically.
Path creation tool	This tool draws Bezier curves in order to create paths. Paths are akin to selections and can be saved for use later on in the image-editing process. Creating a Bezier curve is not too hard to do: just click and drag to draw a curve. Each extra click you make will define a new curve, which will be joined to the last one. To turn the path into a selection, click the button at the bottom of the toolbar.
Color picker	This lets you see the RGB, HSV, or CMYK values of any color within the image. Simply click the mouse within the image.
Zoom tool	Click to zoom into the image, right-click to see various zoom options, and hold down the Alt key while clicking to zoom out.
Measurer	This tool measures distances between two points (in pixels) and also angles. Just click and drag to use it. The measurements will appear at the bottom of the image window.
Move tool	Click and drag to move any selection areas within the image, as well as rearrange the positioning of various layers.
Crop tool	Click and drag to define an area of the image to be cropped. Anything outside the selection area you create will be discarded.
Rotate tool	This tool rotates any selections you make and can also rotate entire layers. It opens a dialog box in which you can set the rotation manually. Alternatively, you can simply click and drag the handles behind the dialog box to rotate by hand.
Scale tool	Known in some other image editors as “transform,” this tool lets you resize the selection area or layer. It presents a dialog box where you can enter numeric values, or you can click and drag the handles to resize by hand.
Shear tool	This tool lets you transform the image by shearing it. Slant a selection by clicking and dragging the corners of the selection area (if the selection area isn’t square, a rectangular grid will be applied to it for the purposes of transformation).
Perspective tool	This tool lets you transform a selection by clicking and dragging its four corners and independently moving them without affecting the other corners. In this way, a sense of perspective can be emulated.
Flip tool	This tool flips a selection or image so that it is reversed on itself, either horizontally (click) or vertically (hold down Ctrl and click).
Text tool	Click on the image to add text.
Fill tool	Fill a particular area with solid color, according to the color selected in the color box below.

Table 20-1. *The GIMP Image-Editing Tools (Continued)*

Tool	Description of Use
Gradient fill	This tool will create a gradient fill based on the foreground and background colors by clicking and dragging.
Pencil tool	This tool lets you draw individual pixels when zoomed in, or hard-edge lines when zoomed out. Simply click and drag to draw freehand, and hold down Shift to draw lines between two points.
Brush tool	This tool lets you draw on the picture in a variety of brush styles to create artistic effects. A brush can also be created from an image, allowing for greater versatility.
Erase tool	Rather like the Brush tool in reverse, this tool deletes whatever is underneath the cursor. If layers are being used, the contents of the layer beneath will become visible.
Airbrush tool	This tool is also rather like the Brush tool, in that it draws on the picture in a variety of styles. However, the density of the color depends on the length of time you press the mouse button. Tap the mouse button, and only a light color will appear. Press and hold the mouse button, and the color will become more saturated.
Ink tool	This tool is like the Brush tool except that, rather like an ink pen, the faster you draw, the thinner the brush stroke is.
Pattern stamp	Commonly known as the clone tool, this is a popular image-editing tool. It is able to copy one part of an image to another via drawing with a brush-like tool. The origin point is defined by holding down Ctrl and clicking.
Blur/sharpen tool	Clicking and drawing on the image will spot blur or sharpen the image, depending on the settings in the tool options area in the lower half of the toolbar.
Smudge tool	As its name suggests, clicking and drawing with this tool will smudge the image, rather like rubbing a still-wet painting with your finger (except slightly more precise).
Burn and dodge tool	This tool lets you spot lighten and darken an image by clicking and drawing on the image. The results depend on the settings in the tool options part of the window.

Directly beneath the image-editing tool icons, on the right, is an icon that shows the foreground and background colors that will be used when drawing with tools such as the Brush. To define a new color, double-click either the foreground (top) or background (bottom) color box. To the left is the pattern selector, which lets you choose which patterns are used with tools such as the Brush.

Beneath these icons, you'll see the various options for the selected tool. By using the buttons at the bottom of the window, you can save the current tool options, load tool options, and delete a previously saved set of tool options. Clicking the button on the bottom right lets you revert to the default settings for the tool currently being used (useful if you tweak too many settings!).

Next to the toolbar window is the Layers dialog box. This can be closed for the moment, although you can make it visible again later, if you wish.

Editing Images with The GIMP

After you've started The GIMP (and assigned it a virtual desktop), you can load an image by selecting File ► Open. The browser dialog box offers a preview facility on the right of the window.

You will probably need to resize the image window so that it fits within the remainder of the screen. You can then use the Zoom tool (see Table 20-1) to ensure that the image fills the editing window, which will make working with it much easier.

You can save any changes you make to an image by right-clicking it and selecting File ► Save As. You can also print the image from the same menu.

Before you begin editing with The GIMP, you need to be aware of some essential concepts that are vital to understand in order to get the most from the program:

Copy, cut, and paste buffers: Unlike Windows programs, The GIMP lets you cut or copy many selections from the image and store them for use later. It refers to these saved selections as *buffers*, and each must be given a name for future reference. A new buffer is created by selecting an area using any of the selection tools, then right-clicking within the selection area and selecting Edit ► Buffer ► Copy Named (or Cut Named). Pasting a buffer back is a matter of right-clicking the image and selecting Edit ► Buffer ► Paste Named.

Paths: The GIMP paths are not necessarily the same as selection areas, although it's nearly always possible to convert a selection into a path and vice versa (right-click within the selection or path and look for the relevant option on the Edit menu). In general, the tools used to create a path allow the creation of complex shapes rather than simple geometric shapes, as with the selection tools. You can also be more intricate in your selections, as shown in the example in Figure 20-2. You can save paths for later use. To view the Paths dialog box, right-click the image and select Dialogs ► Paths.

■ **Tip** Getting rid of a selection or path you've drawn is easy. In the case of a path, simply click on any other tool. This will cause the path to disappear. To get rid of a selection, select any selection tool and quickly click once on the image, being careful not to drag the mouse while doing so.




Figure 20-2. Paths allow for more elaborate and intricate selections, such as those that involve curves.

Layers: In The GIMP (along with most other image-editing programs), layers are like transparent sheets of plastic that are placed on top of the image. Anything can be drawn on each individual transparent sheet, and many layers can be overlaid in order to create a complicated image. Layers also let you cut and paste parts of the image between them. It's also possible to apply effects and transformations to a single layer, rather than to the entire image. The Layers dialog box, shown in Figure 20-3, appears by default, but if you closed it earlier, you can open it again by right-clicking the image and selecting **Dialogs** ► **Layers**. The layers can be reordered by clicking and dragging them in the dialog box. In addition, the blending mode of each layer can be altered. This refers to how it interacts with the layer below it. For example, its opacity can be changed so that it appears semitransparent, thereby showing the contents of the layer beneath.



Figure 20-3. Set the opacity of various layers by clicking and dragging the relevant slider in the Layers dialog box.

Making Color Corrections

The first step when editing most images is to correct the brightness, contrast, and color saturation. This helps overcome some of the deficiencies that are inherent in digital photographs or scanned-in images. To do this, right-click the image and select **Layers > Colors**. You'll find a variety of options to let you tweak the image, allowing you a lot of control over the process.

For trivial brightness and contrast changes, selecting the **Brightness/Contrast** menu option will open a dialog box where clicking and dragging the sliders will alter the image. The changes you make will be previewed on the image itself, so you should be able to get things just right.

Similarly, the **Hue/Saturation** option will let you alter the color balance and also the strength of the colors (the saturation) by clicking and dragging sliders. By selecting the color bar options at the top of the window, you can choose individual colors to boost. Clicking the **Master** button will let you once again alter all colors at the same time.

The trouble with clicking and dragging sliders is that it relies on human intuition. This can easily be clouded by a badly calibrated monitor, which might be set too dark or too light. Because of this, The GIMP offers another handy option: **Levels**.

To access the Levels feature, right-click the image and select Layer ► Colors ► Levels. This presents a chart of the brightness levels in the photo and lets you set the dark, shadows, and highlight points, as shown in Figure 20-4. Three sliders beneath the chart represent, from left to right, the darkest point, the midtones (shadows), and the highlights within the picture. The first step is to set the dark and light sliders at the left and right of the edges of the chart. This will make sure that the range of brightness from the lightest point to the darkest point is set correctly. The next step is to adjust the middle slider so that it's roughly in the middle of the highest peak within the chart. This will accurately set the midtone point, ensuring an even spread of brightness across the image.

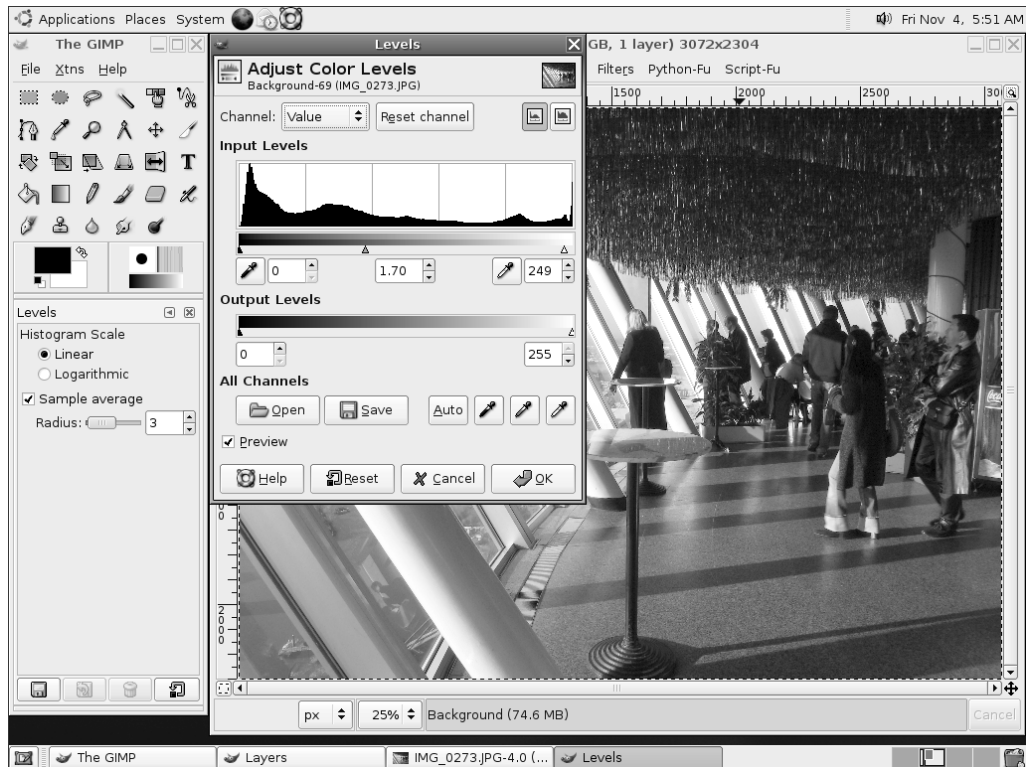


Figure 20-4. The Levels function can be used to accurately set the brightness levels across an image.

A little artistic license is usually allowed at this stage and, depending on the effect on the photo, moving the midtone slider a little to the left and/or right of the highest peak might produce more acceptable results. However, be aware that the monitor might be showing incorrect brightness/color values.

Cropping and Cloning

After you've adjusted the colors, you might want to use the Crop tool (see Table 20-1) to remove any extraneous details outside the focus of the image. For example, in a portrait of someone taken from a distance away, you might choose to crop the photo to show only the person's

head and shoulders, or you might separate a group of people from their surroundings, as shown in Figure 20-5.

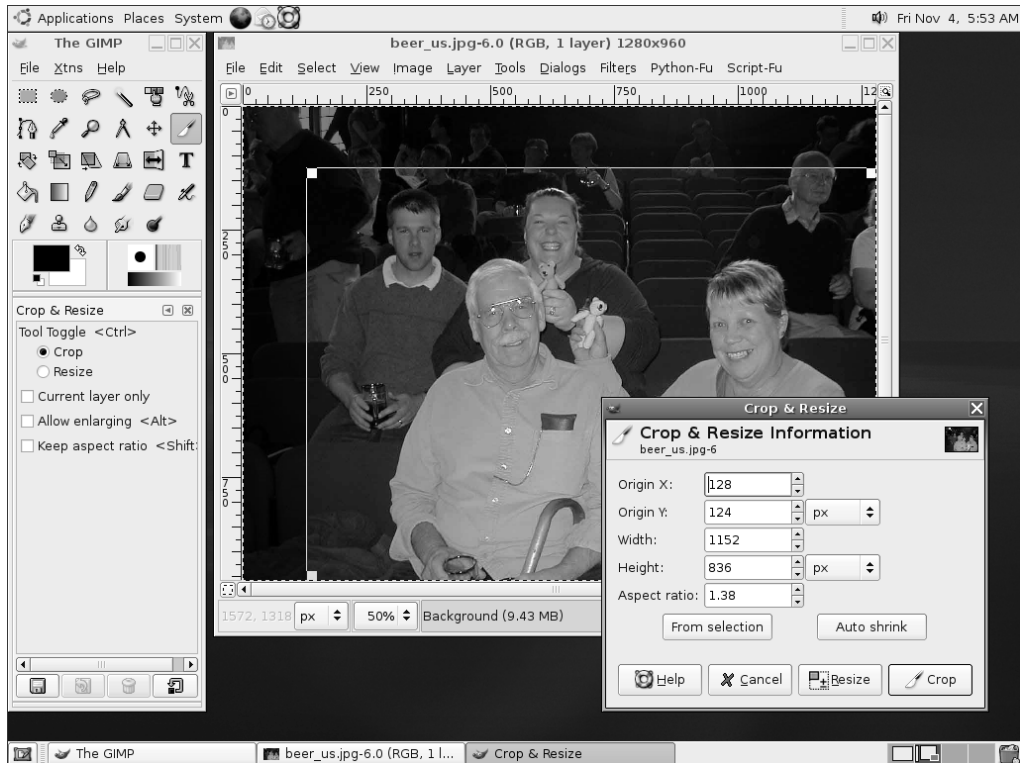


Figure 20-5. You can use the Crop tool to remove any irrelevant details surrounding the subject of your photo.

You might also want to use the Clone tool to remove facial blemishes. Start by using the Zoom tool to close in on the area. If the blemish is small, you might need to zoom in quite substantially. Then try to find an area of skin that is clear and from which you can copy. Hold down Ctrl and click in that area. Then click and draw over the blemish. The crosshair indicates the area from which you're copying.

Sharpening

One final handy trick employed by professional image editors to give their photos a shot in the arm is to use the Sharpen filter. This has the effect of adding definition to the image and negating any slight blur caused by things such as camera shake or poor focusing. To apply the Sharpen filter, right-click the image and select **Filters > Enhance > Sharpen**.

As shown in Figure 20-6, a small preview window will show the effect of the sharpening on the image (you might need to use the scroll bars to move to an appropriate part of the image). Clicking and dragging the slider at the bottom of the dialog box will alter the severity of the

sharpening effect. Too much sharpening can ruin a picture, so be careful. Try to use the effect subtly.

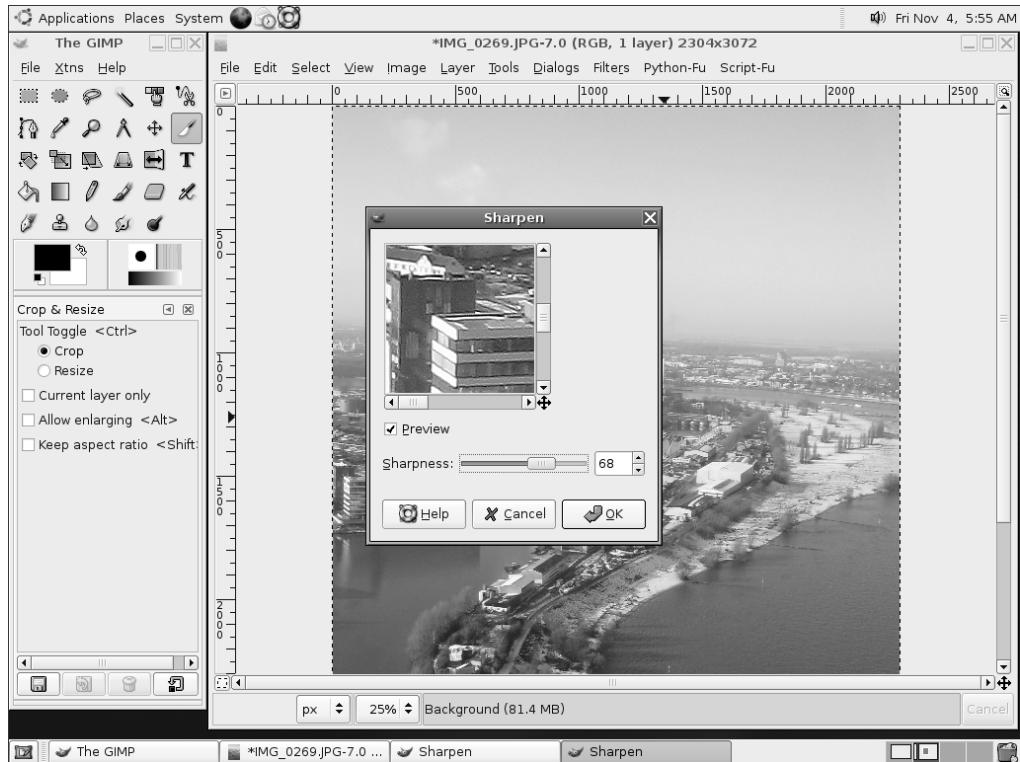


Figure 20-6. Sharpening an image can give it a professional finish by adding definition.

The Sharpen filter is just one of many filters you can apply in The GIMP, as explained in the next section.

Applying Filters

Like other image-editing programs, The GIMP includes many filters to add dramatic effects to your images with little, if any, user input. Filters are applied either to the currently selected layer or to a selection within the layer. To apply a filter, right-click the image and choose the relevant menu option. If you don't like an effect you've applied, you can reverse it by selecting **Edit** ► **Undo**, or by pressing **Ctrl+Z**.

The submenus offer filters grouped by categories, as follows:

Blur: These filters add various kinds of blur to the image or selection. For example, Motion Blur can imitate the effect of photographing an object moving at speed with a slow shutter. Perhaps the most popular blur option is Gaussian Blur, which has the effect of applying a soft and subtle blur.

Color: This option includes many technical filters, mostly of interest to image technicians or those who want to uncover and otherwise manipulate the color breakdown within an image. However, Filter Pack might appeal to the general user. This filter can quickly adjust the hue, saturation, and other values within the image. Also of interest is Colorify, which can tint the image to any user-defined color. Figure 20-7 shows an example of using the Colorify filter.

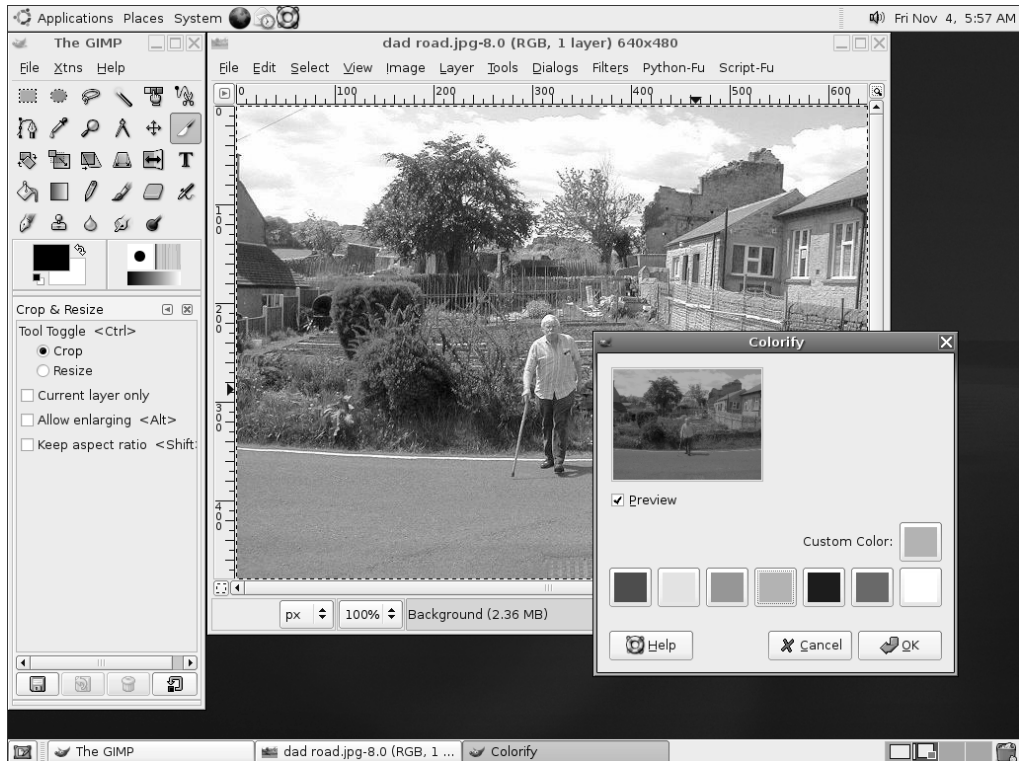


Figure 20-7. *The Colorify filter can be used to add a sepia-like effect to a picture.*

Noise: This collection of filters is designed to add speckles or other types of usually unwanted artifacts to an image. These filters are offered within The GIMP for their potential artistic effects, but they can also be used to create a grainy film effect—simply click Scatter RGB.

Edge Detect: This set of filters can be used to automatically detect and delineate the edges of objects within an image. Although this type of filter can result in some interesting results that might fall into the category of special effects, it's primarily used in conjunction with other tools and effects.

Enhance: The Enhance effects are designed to remove various artifacts from an image or otherwise improve it. For example, the Despeckle effect will attempt to remove unwanted noise within an image (such as flecks of dust in a scanned image). The Sharpen filter discussed in the previous section is located here, as is the Unsharp Mask, which offers a high degree of control over the image-sharpening process.

Generic: In this category, you can find a handful of filters that don't seem to fall into any other category. Of particular interest is the Convolution Matrix option, which lets you create your own filters by inputting numeric values. According to The GIMP's programmers, this is designed primarily for mathematicians, but it can also be used by others to create random special effects. Simply input values and then preview the effect.

Glass Effects: As the name suggests, these filters can apply effects to the image to imitate the effects that come about when glass is used to produce an image. For example, the Apply Lens filter will apply the same kind of distortion caused by various wide-angle lenses used on cameras, as shown in Figure 20-8.



Figure 20-8. *The Glass Effects ► Apply Lens filter can be used to imitate a fish-eye lens.*

Light Effects: Here, you will find filters that imitate the effects that light can have on a picture, such as adding sparkle effects to highlights or imitating lens flare caused by a camera's lens.

Distorts: As the name of this category of filters suggests, the effects here distort the image in various ways. For example, Whirl and Pinch allow you to tug and push the image to distort it (to understand what is meant here, imagine the image is printed on rubber and then pinching or pushing the surface). This category also contains other special effects, such as Pagecurl, which imitates the curl of a page at the bottom of the picture.

Artistic: These filters allow you to add painterly effects to the image, such as making it appear as if the photo has been painted in impressionistic brushstrokes, or as if it were painted on canvas by overlaying the texture of canvas onto the picture. Figure 20-9 shows an example of applying a filter for an oil painting effect.




Figure 20-9. The Artistic effects can be used to give images an oil painting effect.

Map: These filters aim to manipulate the image by treating it like a piece of paper that can be folded in various ways and also stuck onto 3D shapes (a process referred to as *mapping*). Because the image is treated as if it were a piece of paper, it can also be copied, and the copies placed on top of each other to create various effects.

Render: Here, you'll find filters designed to create new images from scratch, such as clouds or flame effects. They obliterate anything that was previously underneath on that particular layer or within that selection, and the original image has no bearing on what is generated by the filter.

Web: Here, you can create an image map for use in a web page. An *image map* is a single image broken up into separate hyperlinked areas, typically used on a web page as a sophisticated menu. For example, an image map is frequently used for a geographical map on which you can click to get more information about different regions.

Animation: These filters aim to manipulate and optimize GIF images, which are commonly used to create simple animated images for use on web sites.

Combine: Here, you'll find filters that combine two or more images into one.

Toys: These are so-called “Easter Eggs,” which aren't designed to manipulate the image, but are present in the program as harmless animations for the user to enjoy. They're created by the programmers of The GIMP as a way of thanking you for using their program.

■ **Tip** The GIMP also includes Script-Fu and Python-Fu, scripting languages that can be used to daisy-chain several commands together to produce a particular effect or to automate a particular image-editing process. For more information, search for “script-fu” or “python-fu” using your favorite web search engine.

Summary

In this chapter, we took a look at image editing under Ubuntu. This has involved an examination of one of the best programs available for the task under any operating system: The GIMP.

You learned how to start The GIMP and about some of the basic principles behind it. Next, we discussed some of the functions contained within The GIMP, including the image filters provided with the program.

In the next part of the book, we move on from multimedia to look at another core component of Ubuntu: the OpenOffice.org suite, which provides word processing, spreadsheet, presentation, and other functions.

PART 6



Office Tasks



Making the Move to OpenOffice.org

You might be willing to believe that you can get a complete operating system for no cost. You might even be able to accept that this offers everything Windows does and much more. But one stumbling block many people have is in believing a Microsoft Office-compatible office suite comes as part of the zero-cost bundle. It's a step too far. Office costs hundreds of dollars—are they expecting us to believe that there's a rival product that is free?

Well, there is, and it's called OpenOffice.org. It comes preinstalled with Ubuntu, as well as most other Linux distributions, making it the Linux office suite of choice. It's compatible with most Microsoft Office files, too, and even looks similar and works in a comparable way, making it easy to learn. What more could you want?

Office Similarities

OpenOffice.org started life as a proprietary product called Star Office. Sun Microsystems bought the company behind the product and released its source code in order to encourage community development. This led to the creation of the OpenOffice.org project, a collaboration between open-source developers and Sun. This project has released several new versions of OpenOffice.org, and at the time of writing, has just released version 2. This is the version supplied with Ubuntu.

Note For what it's worth, Sun still sells Star Office. This is based on the OpenOffice.org code, so it's effectively the same program. However, in addition to the office suite itself, Sun includes several useful extras such as fonts, templates, and the all-important technical support, which you can contact if you get stuck trying to undertake a particular task.

OpenOffice.org features a word processor, spreadsheet program, presentation package, drawing tool (vector graphics), web site creation tool, database program, and several extras.

As such, it matches Microsoft Office almost blow-for-blow in terms of core functionality. See Table 21-1 for a comparison of core packages.

Table 21-1. *How the Office and OpenOffice.org Suites Compare*

Microsoft Office	OpenOffice.org	Function
Word	Writer	Word processor
Excel	Calc	Spreadsheet
PowerPoint	Impress	Presentations
Visio	Draw ¹	Technical drawing/charting
FrontPage	Writer ²	Web site creation
Access	Base ³	Database

¹ *Draw is a vector graphics creation tool akin to Adobe Illustrator. Creating flow charts or organizational diagrams is one of many things it can do.*

² *Writer is used for word processing and HTML creation; when switched to Web mode, its functionality is altered appropriately.*

³ *Writer and Calc can be coupled to a third-party database application such as MySQL or Firebird; however, Office.org also comes with the Base relational database.*

You should find the functionality within the packages is duplicated, too, although some of the very specific features of Microsoft Office are not in OpenOffice.org. But OpenOffice.org also has its own range of such tools not yet found in Microsoft Office!

OpenOffice.org does have a couple of notable omissions. Perhaps the main one is that it doesn't offer a directly comparable Outlook replacement. However, as discussed in Chapter 27, the Evolution application offers an accurate reproduction of Outlook, with e-mail, contacts management, and calendar functions all in one location. In Ubuntu, you'll find Evolution on the Applications ► Internet menu. Evolution isn't directly linked to OpenOffice.org, but it retains the overall Ubuntu look, feel, and way of operating.

OpenOffice.Org Key Features

Key features of OpenOffice.org, apart from the duplication of much of what you find in Microsoft Office, include the ability to export documents in Portable Document Format (PDF) format across the entire suite of programs. PDF files can then be read on any computer equipped with PDF display software, such as Adobe Acrobat Reader.

In addition, OpenOffice.org features powerful accessibility features that can, for example, help those with vision disabilities use the programs more effectively. For those who are technically minded, OpenOffice.org can be extended very easily with a variety of plug-ins, which allow the easy creation of add-ons using many different programming languages.

Although OpenOffice.org largely mirrors the look and feel of Microsoft Office, it adds its own flourishes here and there. This can mean that some functions are located on different menus, for example. However, none of this poses a challenge for most users, and OpenOffice.org is generally regarded as very easy to learn.

File Compatibility

As well as core feature compatibility, OpenOffice.org is also able to read files from Microsoft Office versions up to and including Office 2003, the latest version of Office at the time of writing. When a new version of Office comes out with a new file format, future versions of OpenOffice.org will support it. This is just one more reason why you should regularly update Ubuntu online in order to make sure you're running the very latest versions of each program.

Note It's fair to say that many people still use the older Office file formats, even if they're using the latest version of Office. This is done to retain compatibility with other users who may not yet have upgraded.

Although file compatibility problems are rare, two issues occasionally crop up when opening Microsoft Office files in OpenOffice.org:

VBA compatibility: OpenOffice.org isn't compatible with Microsoft Office Visual Basic for Applications (VBA). It uses a similar but incompatible internal programming language. This means that Microsoft macros within a document probably won't work when the file is imported to OpenOffice.org. Such macros are typically used in Excel spreadsheets designed to calculate timesheets, for example. In general, however, only high-end users use VBA.

Document protection: OpenOffice.org is unable to open any Office files that have a password, either to protect the document from changes or to protect it from being viewed. Theoretically, it would be easy for OpenOffice.org's programmers to include such functionality, but the laws of many countries make creating such a program feature illegal (it would be seen as a device to overcome copy protection). The easiest solution is to ask whoever sent you the file to remove the password protection. For what it's worth, OpenOffice.org's has its own form of password protection.

If you find that OpenOffice.org isn't able to open an Office file saved by your colleagues, you can always suggest that they, too, make the switch to OpenOffice.org. They don't need to be running Ubuntu to do so. Versions are available to run on all Windows platforms, as well as on the Apple Macintosh. As with the Ubuntu version, they're entirely free of charge. Indeed, for many people who are running versions of Office they've installed from "borrowed" CDs, OpenOffice.org offers a way to come clean and avoid pirating software. For more details and to download OpenOffice.org, visit www.openoffice.org.

Once your colleagues have made the switch, you can exchange files using OpenOffice.org's native format, or opt to save files in the Office file format. Figure 21-1 shows the file type options available in OpenOffice.org's word processor component's Save As dialog box.

Note OpenOffice.org also supports Rich Text Format (RTF) text documents and comma-separated value (CSV) data files, which are supported by practically every office suite program ever made.

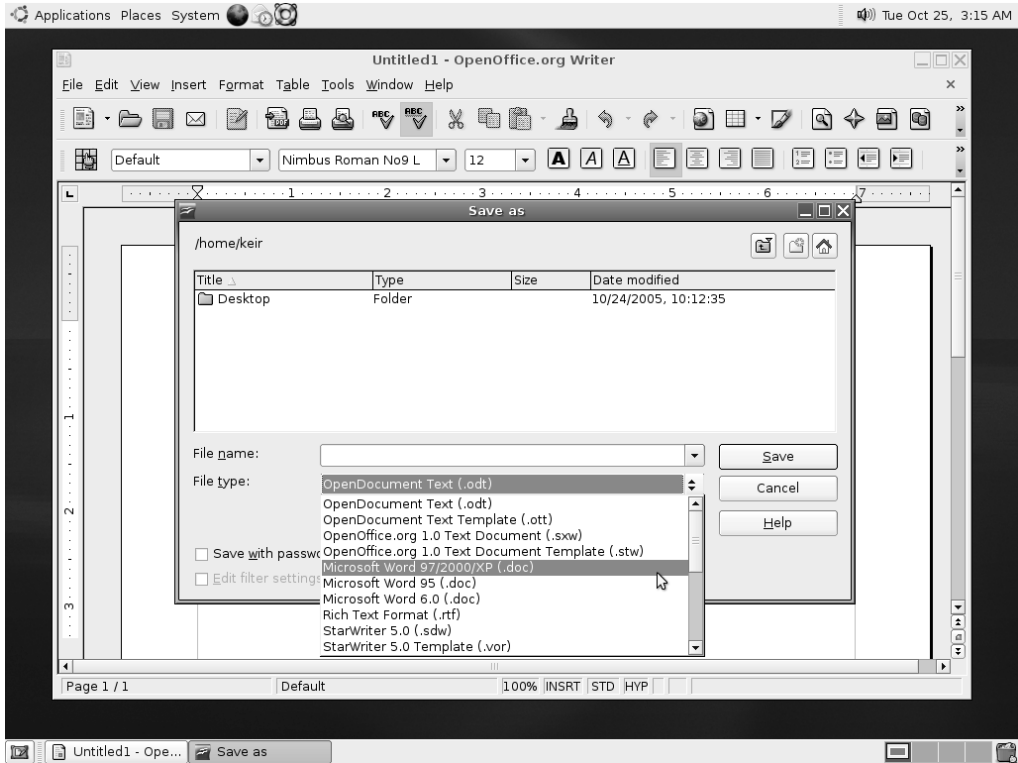


Figure 21-1. All the OpenOffice.org components are fully compatible with Microsoft Office file formats.

When it comes to sharing files, there's another option: save your files in a non-Office format such as PDF or HTML. OpenOffice.org is able to export documents in both formats, and most modern PCs equipped with Adobe Acrobat or a simple web browser will be able to read them. However, while OpenOffice.org can open and edit HTML files, it can export documents only as PDF files, so this format is best reserved for files not intended for further editing.

OPEN DOCUMENT FORMAT

One of the principles behind all open-source software is the idea of open file formats. This means that if someone creates a new open-source word processor, he also makes sure that the technology behind the file format is explained, so that other people can adapt their programs to read and/or save in that file format.

Compare this to Microsoft's approach, which is to not only attempt to keep its file formats secret, but also to patent the formats so that no one else can use them without permission. The only way for Office competitors like OpenOffice.org to understand Office file formats is to reverse-engineer them—to try to figure out how they work by taking them apart. This is a time-consuming process and, increasingly, resides in a legal gray area in many countries.

To meet the goals of open-source software, the OpenDocument Format (ODF) was created. This is a completely open and free to use office document file format that all software suites can adopt. The idea is that ODF will make swapping files between all office suites easy.

Sadly, Microsoft has decided not to support ODF and is sticking with its own proprietary file formats. However, Massachusetts has recently adopted ODF, as have several other local governments in countries all around the world. There's little doubt that it won't be long before ODF will become one of the main ways of disseminating and sharing documents online.

Perhaps it goes without saying that OpenOffice.org fully supports ODF. In fact, ODF was originally based on OpenOffice.org file formats.

The Right Fonts

One key to compatibility with the majority of Microsoft Office files is ensuring you have the correct fonts. This is an issue even when using Windows. It's very common to open an Office document to find the formatting incorrect because you don't have the fonts used in the construction of the document.

Although most Windows systems have many fonts, most people tend to rely on a handful of core fonts, which are default on most Windows installations: Arial, Tahoma, Verdana, Trebuchet MS, and Times New Roman (MS Comic Sans might also be included in that list, although it isn't often used within business documents).

You can obtain these fonts and install them on your Ubuntu system in several ways. Here, we'll cover two methods: copying your fonts from Windows and installing Microsoft's TrueType Core Fonts. The latter method is by far the easier way of undertaking this task.

Copying Windows Fonts

If you dual-boot Ubuntu with Windows, you can delve into your Windows partition's font folder and copy across every font you have available under Windows. This involves not only copying the files across, but also rebuilding several system files, which adds an extra level of complication. This method is useful if you wish to copy across *all* the fonts you use under Windows. If you wish to get just Arial and Times New Roman, you might want to skip ahead to the next section.

Caution Installing Windows fonts under Ubuntu is a legally gray area. Technically speaking, there's no reason why you shouldn't be able to use the fonts under Ubuntu. Purchasing Windows as well as any software running on it should also have meant you purchased a license to use the fonts. But the license document for Windows XP makes no mention of font licensing, so the matter is far from clear-cut.

To copy your fonts, click Applications ► Accessories ► Terminal and, in the window that appears, type the following series of commands:

```
sudo mount /dev/hda1 /mnt/ -o umask=0222
sudo cp /mnt/WINDOWS/Fonts/*.ttf /usr/share/fonts
sudo chown root.root /usr/share/fonts/*.ttf
sudo mkfontdir /usr/share/fonts/
cd /
sudo umount /mnt/
```

In the second line:

```
sudo cp /mnt/WINDOWS/Fonts/*.ttf /usr/share/fonts
```

you may need to substitute the path to the Windows font directory, taking into account lower-case and uppercase letters. This path was relevant for a Windows XP Home machine; other versions of Windows, such as Windows 98 or Me, may use different paths. For example, on a machine with Windows 98 installed, the path was `/mnt/windows/fonts/`.

These commands mount the Windows partition, copy the fonts to your Ubuntu partition, and then rebuild the system file that catalogs the fonts.

After you've issued the commands, close all program windows, click **System ► Log Out**, and opt to log out of the system. Then log back in again. You should now find your Windows fonts are available in all applications, including OpenOffice.org.

Installing TrueType Core Fonts

If you don't want to undertake the font-copying maneuver, you can download and install Microsoft's TrueType Core Fonts. This package contains common Windows fonts, including Arial and Times New Roman.

The first few steps in the following procedure take you through setting up Ubuntu's software download feature to accept software from the Ubuntu online repositories, which is covered in more detail in Chapter 8. If you've already configured the Synaptic Package Manager to access these repositories, open the Synaptic Package Manager and start with step 5.

1. Click **System ► Administration ► Synaptic Package Manager**. Enter your password to continue.
2. In the program window that appears, click **Settings ► Repositories**. Then click the **Settings** button. In the dialog box that appears, put a check next to **Show Disabled Software Sources**.
3. The list is now lengthened by several new entries. Look for the second entry entitled **Ubuntu 5.10 "Breezy Badger" (Binary)** and make sure there is a check in the box alongside it. Then click the **Edit** button.

4. In the line headed Sections, type **multiverse** and **universe** after the two current entries. This line should read “main restricted multiverse universe.” Click OK, and then click OK again to return to the main program window. Opt to reload new package information when asked.
5. Click the Search icon and type `msttcorefonts`. Click the package’s leftmost icon when it appears in the results list, and then select Mark for Installation, as shown in Figure 21-2. You’ll be warned an additional program needs to be installed; this is fine. Then click Apply to install the fonts.
6. Close all program windows, click System ► Log Out, and opt to log out of the system. Then log back in again. You should now find your Windows fonts are available in all applications, including OpenOffice.org.

UPDATING OPENOFFICE.ORG

Although the version of OpenOffice.org supplied with Ubuntu purports to be version 2.0, it’s actually version 1.9.129. This is a preview release that was included to meet Ubuntu’s release schedule (at the time, version 2.0 wasn’t available). However, upgrading to the full release of 2.0 is easy.

1. Select System ► Administration ► Synaptic Package Manager.
2. Click Settings ► Repositories, and then click the Add button.
3. Click the Custom button. In the APT Line box of the dialog box, enter the following (note the space before the `./` at the end of the line):

```
http://people.ubuntu.com/~doko/OOo2 ./
```
4. Click the Add Repository button.
5. Click OK in the Software Preferences dialog box. Then click Yes to reload the package list.
6. Close the Synaptic Package Manager.
7. After a few seconds, the software update icon should appear in the notification area. Click this icon.
8. You’ll see that a new version of OpenOffice.org is available. Click the Install button to download and install it. Be aware that OpenOffice.org weighs in at around 80MB, so will take some time to download.

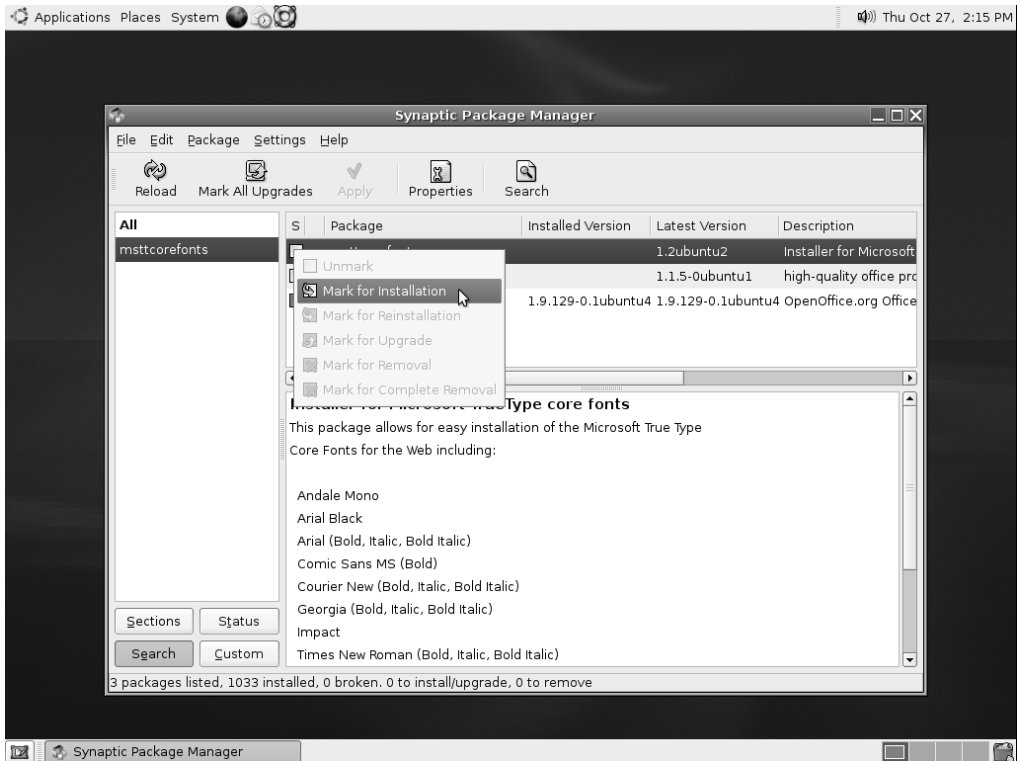


Figure 21-2. Vital Microsoft fonts are just a download away courtesy of the Synaptic Package Manager.

OTHER LINUX OFFICE SUITES

OpenOffice.org is widely regarded as one of the best Linux office suites, but it's not the only one. Its main open-source competitor is KOffice. KOffice tightly integrates into the KDE desktop and mirrors much of its look and feel. It includes a word processor, spreadsheet, presentation package, flow-charting tool, database-access tool, and much more. As with OpenOffice.org, in most cases, you can load and save Microsoft Office files. For more details, see its home page at www.koffice.org. It's available with Ubuntu, too. Just use the Synaptic Package Manager to search for and install it.

In addition, there are several open-source office applications that aren't complete office suites. For example, AbiWord is considered an excellent word processor, which packs in a lot of features but keeps the user interface very simple. It's partnered by Gnumeric, a spreadsheet application that is developed separately (although both aim to be integrated into the GNOME desktop environment). For more details, see www.abiword.com and www.gnome.org/projects/gnumeric/, respectively. You can also find both of these programs in the Ubuntu software repositories (use the Synaptic Package Manager to search for them).

If you don't mind paying for proprietary software, Hancom Office is extremely popular in Asian countries, and an English language version is also available. It offers a word processor, spreadsheet, presentations package, and more. A 30-day trial version and more details are available from <http://en.hancom.com>.

Summary

This chapter was a general introduction to OpenOffice.org, providing an overview of what you can expect from the programs within the suite. In particular, we focused on the extent of the suite's similarities with Microsoft Office and discussed issues surrounding file compatibility with Microsoft Office. We also looked at how Windows fonts can be brought into Ubuntu, which aids in successfully importing and creating compatible documents.

In the next chapter, you'll learn about the configuration options globally applicable to the suite, as well as common functions provided across all the programs.



OpenOffice.org Overview

All the programs in the OpenOffice.org suite rely on a common interface, and therefore look and operate in a similar way. They are also configured in an identical way, and all rely on central concepts such as wizards, which guide you through the creation of particular types of documents. In addition, many components within the suite are shared across the various programs. For example, the automatic chart creation tool within Calc can also be used within Writer.

In this chapter, we'll look at the OpenOffice.org suite as a whole, and explain how it's used and configured. In the following chapters, we'll examine some specific programs in the suite.

Introducing the Interface

If you've ever used an office suite, such as Microsoft Office, you shouldn't find it too hard to get around in OpenOffice.org. As with Microsoft Office, OpenOffice.org relies primarily on toolbars, a main menu, and separate context-sensitive menus that appear when you right-click. In addition, OpenOffice.org provides floating palettes that offer quick access to useful functions, such as paragraph styles within Writer.

Figure 22-1 provides a quick guide to the OpenOffice.org interface, showing the following components:

- **Menu bar:** The menus provide access to most of the OpenOffice.org functions.
- **Standard toolbar:** This toolbar provides quick access to global operations, such as saving, opening, and printing files, as well as key functions within the program being used. The Standard toolbar appears in all OpenOffice.org programs and also provides a way to activate the various floating palettes, such as the Navigator, which lets you easily move around various elements within the document.
- **Formatting toolbar:** As its name suggests, this toolbar offers quick access to text-formatting functions, similar to the type of toolbar used in Microsoft Office applications. Clicking the B icon will boldface any selected text, for example. This toolbar appears in Calc, Writer, and Impress.
- **Ruler:** The ruler lets you set tabs and alter margins and indents (within programs that use rulers).

- **Status bar:** The status bar shows various aspects of the configuration, such as whether Insert or Overtyping mode is in use.
- **Document area:** This is the main editing area.

Most of the programs rely on the Standard and Formatting toolbars to provide access to their functions, and some programs have additional toolbars. For example, applications such as Impress (a presentation program) and Draw (for drawing vector graphics) have the Drawing toolbar, which provides quick access to tools for drawing shapes, adding lines, and creating fills (the blocks of color within shapes).

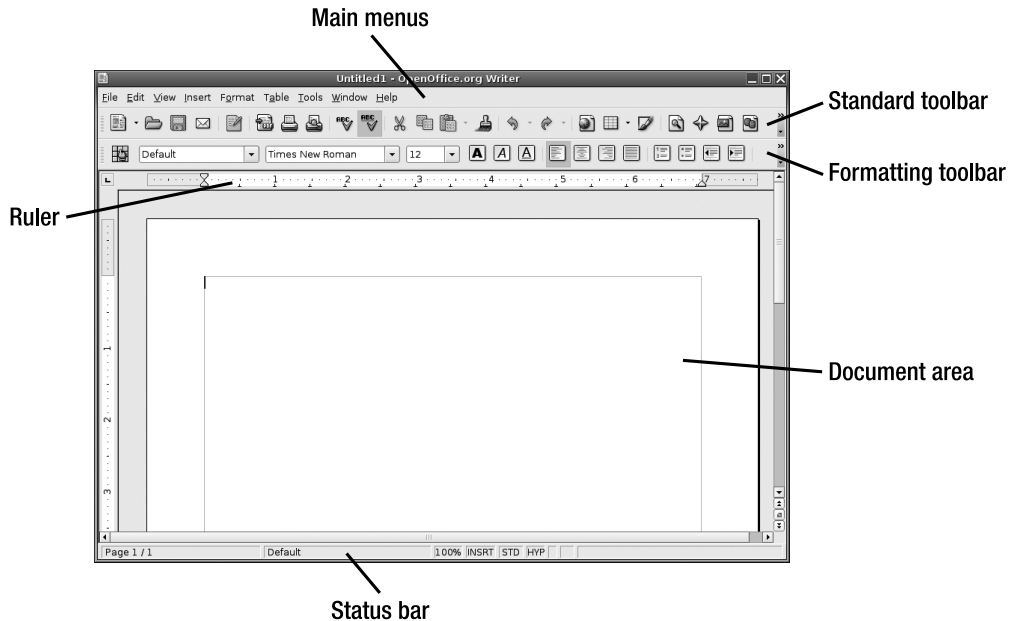


Figure 22-1. *The OpenOffice.org interface has several components.*

Customizing the Interface

You can select which toolbars are visible on your screen, as well as customize those that are already there. You can also add new toolbars and customize the OpenOffice.org menus.

Adding Functions to Toolbars

The quickest way to add icons and functions to any toolbar is to click the two small arrows at the right of a toolbar and select the Visible Buttons entry on the menu that appears. This will present a list of currently visible icons and functions, along with those that might prove useful on that toolbar but are currently hidden. Any option already visible will have a check next to it.

Additionally, you can add practically any function to a toolbar, including the options from the main menus and many more than those that are ordinarily visible. Here are the steps:

1. Click the two small arrows to the right of a toolbar and select the Customize Toolbar option.
2. In the Customize dialog box, click the Add button in the Toolbar Contents section to open the Add Commands dialog box, as shown in Figure 22-2.

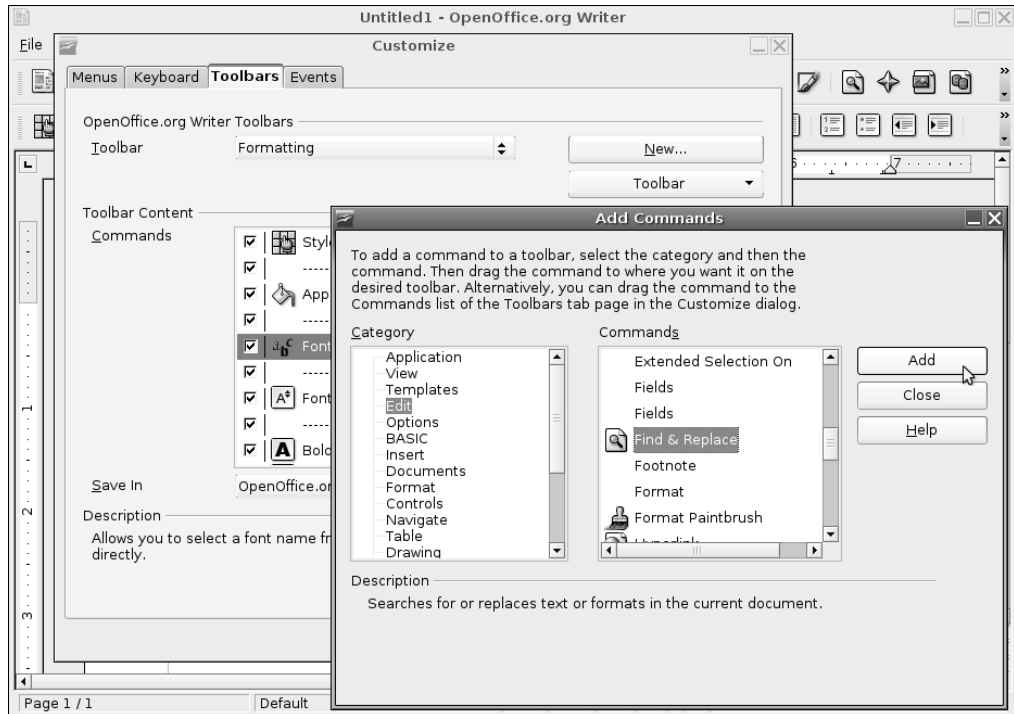


Figure 22-2. Adding a new function to the toolbar is very easy within OpenOffice.org.

3. Choose a category from the list on the left to see the available commands in the list on the right. The categories of functions are extremely comprehensive. For example, under the Format category, you'll find entries related to specific functions, such as shrinking font sizes or setting a shadow effect behind text. Table 22-1 provides brief descriptions of each of the categories listed in the Add Commands dialog box.
4. Select the function you want to add on the right side of the Add Commands dialog box, and then click the Add button.
5. Click the Close button. You'll then see your new function in the list of icons in the Customize dialog box.
6. Click and drag to move the new function left or right on the toolbar itself (you'll see the toolbar itself update when you release the mouse button). Alternatively, you can highlight the icon and click the up and down arrows next to the list.

Table 22-1. *OpenOffice.org Toolbar Customization Categories*

Category	Description
Application	These options relate to the specific OpenOffice.org application you're using. For example, if you select to customize a toolbar within Writer, the Application category menu will offer functions to start AutoPilots (effectively wizards) that will build word processor documents.
View	This category offers options related to the look and feel of the suite, such as which items are visible within the program interface.
Templates	In this category, you'll find options related to the creation and use of document templates.
Edit	This category contains options related to cutting, pasting, and copying items within the document, as well as updating elements within it.
Options	These are various options that relate to configuration choices in OpenOffice.org, allowing you to control how it works.
BASIC	Options under this category relate to the creation and playback of OpenOffice.org macros.
Insert	This category includes options related to inserting objects, such as sound, graphics, and elements from other OpenOffice.org documents.
Documents	This category provides options specific to document control, such as those related to exporting documents as PDF files or simply saving files.
Format	Here, you'll find a range of options related largely to text formatting, but also some concerned with formatting other elements, such as drawings and images.
Controls	Under this heading, you'll find widgets that can be used in conjunction with formulas or macros, such as check boxes, buttons, text box creation tools, and so on.
Navigate	This category offers tools that let you move around a document quickly, such as the ability to quickly edit headers and footers, or move from the top of the page to the end very quickly.
Table	Here, you'll find options related to the creation of tables.
Drawing	Here, you'll find tools related to drawing objects, such as shapes and lines, and also tools for creating floating text boxes.
Graphic	This category presents a handful of options related to manipulating bitmap graphics that are inserted into the document.
Data	Here, you'll find a couple of options related to working with information sources, such as databases.
Frame	These options relate to any frames inserted into the document, such as how elements within the frame are aligned and how text is wrapped around the frame.
Numbering	These are various options related to creating automatic numbered or bulleted lists.
Modify	These options relate to the drawing components within OpenOffice.org and let you manipulate images or drawings in various ways by applying filters.
OpenOffice.org BASIC Macros	Here, you can select from various ready-made macros, which provide some of OpenOffice.org's functions.

Many functions that can be added are automatically given a relevant toolbar icon, but you can choose another icon for a function by selecting the icon in the list in the Customize dialog box, clicking Modify, and then selecting Change Icon. You can also use this method to change an icon that already appears on a toolbar.

Note To delete an icon from a toolbar, click the two small arrows to the right of a toolbar and select the Customize Toolbar option. Select the icon you want to remove, click the Modify button, and choose to delete it.

Adding a New Toolbar

If you want to add your own new toolbar to offer particular functions, you'll find it easy to do. Here are the steps:

1. Click the two small arrows to the right of any toolbar and select Customize Toolbar from the list of options. Don't worry—you're not actually going to customize that particular toolbar!
2. In the Customize dialog box, click the New button at the top right.
3. Give the toolbar a name. The default entry for the Save In field is correct, so you don't need to alter it.
4. Populate the new toolbar, following the instructions in the previous section.
5. Once you've finished, click the OK button.

You should see your new toolbar beneath the main toolbars. To hide it in future, click View ► Toolbars, and then remove the check alongside the name of your toolbar.

Customizing Menus

You can also customize the OpenOffice.org menus. Here are the steps:

1. Select Tools ► Customize from the menu bar.
2. In the Customize dialog box, select the Menu tab.
3. Choose which menu you wish to customize from the Menu drop-down list.
4. Select the position where you wish the new function to appear on the menu, by selecting an entry on the menu function list, and then click the Add button.
5. Add commands to the menu, as described earlier in the "Adding Functions to Toolbars" section.

The up and down arrows in the Customize dialog box allow you to alter the position of entries on the menu. You could move those items you use frequently to the top of the menu, for example.

You can remove an existing menu item by highlighting it in the Customize dialog box, clicking the Modify button, and then clicking Delete.

If you make a mistake, simply click the Reset button at the bottom right of the Customize dialog box to return the menus to their default state.

Configuring OpenOffice.org Options

In addition to the wealth of customization options, OpenOffice.org offers a range of configuration options that allow you to make it work exactly how you wish. Within an OpenOffice.org program select Tools ► Options from the menu to open the Options dialog box, as shown in Figure 22-3.

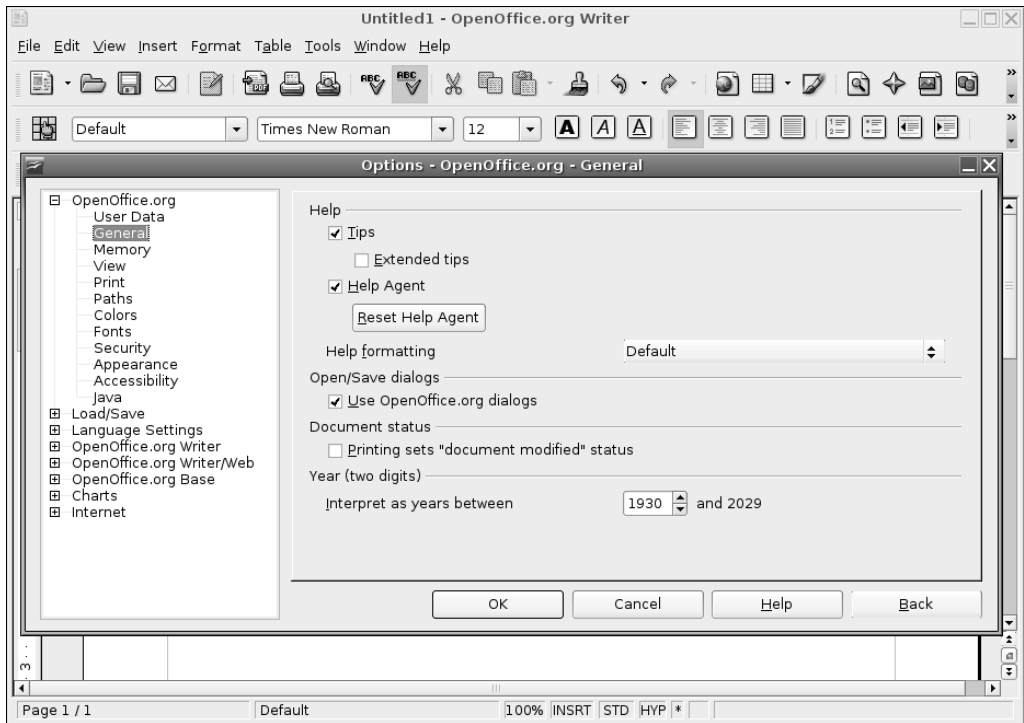


Figure 22-3. OpenOffice.org's main configuration options are accessed by selecting Tools ► Options.

Most of the configuration options offered within each program apply across the suite, but those under the heading of the program's name apply only to the program in use. In other words, to set the options for Writer, you need to use the Writer Options dialog box. But to set global options for the entire suite, you can use any program's Options dialog box.

A variety of options are offered, allowing you to tweak everything from the default file format to the colors used by default within the software. Table 22-2 briefly describes each of the OpenOffice.org configuration options.

Table 22-2. *OpenOffice.org Configuration Options*

Option	Description
OpenOffice.org	
User Data	This is the personal data that will be added to the documents you create. You can leave this area blank if you wish.
General	This offers a handful of miscellaneous options, such as how to handle two-digit dates, when the help system should step in to offer tips, how the help system should be formatted (such as in high resolution for people with vision problems), and whether printing a document is interpreted by OpenOffice.org as modifying it.
Memory	This entry relates to how much system memory OpenOffice.org can use. You can limit the number of undo steps, for example, and alter the cache memory used for holding graphical objects.
View	Here, you can alter the look, feel, and operation of OpenOffice.org. You can define whether the middle mouse button performs a paste operation (which is consistent with how Ubuntu works), or whether it should perform a scrolling function, as with Windows. You can also alter elements such as whether icons appear in menus and if fonts are previewed in the toolbar menu.
Print	This option lets you adjust how printing is handled within OpenOffice.org. The functions relate to those that can stop documents from printing incorrectly, such as reducing any transparency effects within the documents so on-page elements don't appear faint or completely disappear in the final output. (Note that specific print functions are handled within the Print dialog box when you actually print a document.)
Paths	This is where the file paths for user-configured and vital system tools are handled. Generally, there's little reason to edit this list, although you might choose to alter the default location where your documents are saved (simply double-click the My Documents entry to do this).
Colors	Here, you can define the default color palette that appears in the various programs in the suite.
Fonts	By creating entries here, you can automatically substitute fonts within documents you open for others on your system. If you don't have the Microsoft core fonts installed, this might prove useful. For example, you might choose to substitute Arial, commonly used in Microsoft Office documents, for Luxi Sans, one of the sans serif fonts used under Ubuntu.
Security	This option controls which types of functions can be run within OpenOffice.org. For example, you can choose whether macros created by third parties should be run when you open a new document.
Appearance	Here, you can alter the color scheme used within OpenOffice.org, in a similar way to how you can alter the default Ubuntu desktop color scheme. Individual elements within documents and pages can be modified, too.
Accessibility	This option relates to features that might help people with vision disabilities to use OpenOffice.org. For example, you can define whether animated graphics are shown on the screen.
Java	This option lets you control whether you use the Java Runtime Environment, which may be necessary to use some of OpenOffice.org's features.

Table 22-2. *OpenOffice.org Configuration Options (Continued)*

Option	Description
Load/Save	
General	Options here relate to how files are saved. You can select whether the default is to save in OpenOffice.org or Microsoft Office format. Choosing the latter is useful if you share a lot of documents with colleagues who are not running OpenOffice.org.
VBA Properties	This option relates to how Visual Basic for Applications (VBA) code is handled when Microsoft Office documents are opened. Specifically, it ensures that the code isn't lost when the file is saved again.
Microsoft Office	This option provides functions specifically needed to convert or open Microsoft Office files within OpenOffice.org.
HTML Compatibility	Here, you can set options that affect the compatibility of HTML files saved within OpenOffice.org.
Language Settings	
Languages	Here, you can set your local language so that documents are spell-checked correctly. In addition, Asian language support can be activated, which allows for more complex document layout options.
Writing Aids	Under this option, you can activate or deactivate various plug-ins designed to help format documents, such as the hyphenator or the spell-checking component. In addition, you can alter how the spell-checker works, such as whether it ignores capitalized words.
OpenOffice.org Writer	
General	Here, you can alter various options related to the editing of word processor documents, such as which measurements are used on the ruler (centimeter, inches, picas, and so on).
View	Under this option, you can configure the look and feel of the Writer program, such as which scroll bars are visible by default. You can also turn off the display of various page elements, such as tables and graphics.
Formatting Aids	This option lets you choose which symbols appear for "invisible" elements (such as the carriage return symbol or a dot symbol to indicate where spaces have been inserted) in Writer.
Grid	This controls whether page elements will snap to an invisible grid. You can also define the dimensions and spacing of the grid cells here.
Basic Fonts (Western)	This controls which fonts are used by default in the various text styles, such as for the default text and within lists.
Print	This option offers control over printing options specific to Writer, such as which page elements are printed (you might choose to turn off the printing of graphics, for example).
Tables	Here, you can control how tables are created and how you interact with them within Writer. For example, you can control what happens when a table is resized, such as whether the entire table responds to the changes or merely the cell you're resizing.

Table 22-2. *OpenOffice.org Configuration Options (Continued)*

Option	Description
Changes	This option lets you define how changes are displayed when the track changes function is activated.
Compatibility	Here, you can set specifics of how Writer handles the import and export of Microsoft Word documents.
AutoCaption	This offers settings for the AutoCaption feature within Writer.
Mail Merge E-Mail	This option lets you control the sending of e-mail mail merge messages.
OpenOffice.org Writer/Web	
View	Here, you can control the HTML editor component of OpenOffice.org (effectively an extension of Writer). You can control the look and feel of the HTML editor, including which elements are displayed on the screen.
Formatting Aids	As with the similar entry for Writer under Text Document, this option lets you view symbols in place of usually hidden text elements.
Grid	This lets you define a grid that on-screen elements are able to “snap to” in order to aid accurate positioning.
Print	Here, you can define how HTML documents created within OpenOffice.org are printed.
Table	Similar to the Tables entry under Text Document, this controls how tables are created and handled within HTML documents.
Background	This lets you set the default background color for HTML documents.
OpenOffice.org Calc	
General	Here, you can modify miscellaneous options related to Calc, such as which measurement units are used within the program and how the formatting of cells is changed when new data is input.
View	This option relates to the look and feel of Calc, such as the color of the grid lines between cells and which elements are displayed on the screen. For example, you can configure whether zero values are displayed, and whether overflow text within cells is shown or simply truncated at the cell boundary.
Calculate	This option relates to how numbers are handled during certain types of formula calculations, such as those involving dates.
Sort Lists	This option lets you create lists that are applied to relevant cells when the user chooses to sort them. Several lists are predefined to correctly sort days of the week or months of the year.
Changes	This option relates to the on-screen formatting for changes when the track changes function is activated.
Grid	This option lets you configure an invisible grid that stretches across the sheet, and also which page objects can be set to snap to the grid for correct alignment.
Print	This option relates to printing specifically from Calc, such as whether Calc should avoid printing empty pages that might occur within documents.

Table 22-2. *OpenOffice.org Configuration Options (Continued)*

Option	Description
OpenOffice.org Impress	
General	This option refers to miscellaneous settings within the Impress program, such as whether the program should always start with a wizard and which units of measurement should be used.
View	This option relates to the look and feel of Impress, and, in particular, whether certain on-screen elements are displayed.
Grid	This controls whether an invisible grid is applied to the page and whether objects should snap to it.
Print	This option controls how printing is handled within Impress and, in particular, how items in the document will appear on the printed page.
OpenOffice.org Draw	
General	This option relates to miscellaneous settings within Draw (the vector graphics component of OpenOffice.org).
View	Here, you can set specific preferences with regard to which objects are visible on the screen while you're editing with Draw.
Grid	This option relates to the invisible grid that can be applied to the page.
Print	This option lets you define which on-screen elements are printed and which are not printed.
OpenOffice.org Base	
Connections	This option lets you control how any data sources you attach to are handled.
Databases	Here, you can configure which databases are registered for use within Base.
Chart	
Default Colors	Here, you can set the default color palette that should be used when creating charts, usually within the Calc program.
Internet	
Proxy	Here, you can configure network proxy settings specifically for OpenOffice.org, if necessary.
Search	Certain functions within various OpenOffice.org programs let you search the Internet. Here, you can configure how these search functions work.
E-mail	This option lets you specify which program you wish OpenOffice.org to use for e-mail.
Mozilla Plug-in	This function allows integration of OpenOffice.org into the Mozilla and/or Netscape browsers, to allow the viewing of OpenOffice.org documents within the browser window.

Using OpenOffice.org Core Functions

Although the various programs within OpenOffice.org are designed for very specific tasks, they all share several core functions that work in broadly similar ways. In addition, each program is able to borrow components from other programs in the suite.

Using Wizards

One of the core functions you'll find most useful when you're creating new documents is the wizard system, which you can access from the File menu. A wizard guides you through creating a new document by answering questions and following a wizard-based interface. This replaces the template-based approach within Microsoft Office, although it's worth noting that OpenOffice.org is still able to use templates.

A wizard will usually offer a variety of document styles. Some wizards will even prompt you to fill in salient details, which they will then insert into your document in the relevant areas.

Getting Help

OpenOffice.org employs a comprehensive help system, complete with automatic context-sensitive help, called the Help Agent, which will appear if the program detects you're performing a particular task. Usually, the Help Agent takes the form of a light bulb graphic, which will appear at the bottom-right corner of the screen. If you ignore the Help Agent, it will disappear within a few seconds. Clicking it causes a help window to open. Alternatively, you can access the main searchable help file by clicking the relevant menu entry.

Inserting Objects with Object Linking and Embedding

All the OpenOffice.org programs are able to make use of Object Linking and Embedding (OLE). This effectively means that one OpenOffice.org document can be inserted into another. For example, you might choose to insert a Calc spreadsheet into a Writer document.

The main benefit of using OLE over simply copying and pasting the data is that the OLE item (referred to as an *object*) will be updated whenever the original document is revised. In this way, you can prepare a report featuring a spreadsheet full of figures, for example, and not need to worry about updating the report when the figures change. Figure 22-4 shows an example of a spreadsheet from Calc inserted into a Writer document.

Whenever you click inside the OLE object, the user interface will change so that you can access functions specific to that object. For example, if you had inserted an Impress object into a Calc document, clicking within the object would cause the Calc interface to temporarily turn into that of Impress. Clicking outside the OLE object would restore the interface back to Calc.

You can explore OLE objects by selecting **Insert** ► **Object** ► **OLE Object**. This option lets you create and insert a new OLE object, as well as add one based on an existing file.

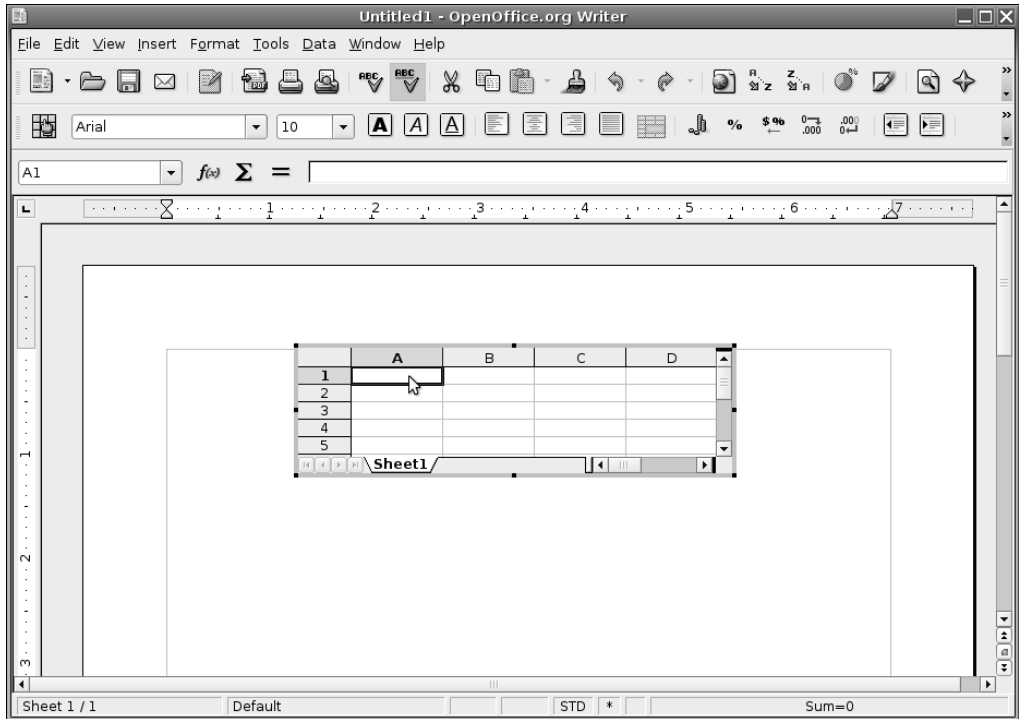


Figure 22-4. Object Linking and Embedding (OLE) lets you incorporate one OpenOffice.org document into another.

Creating Macros

OpenOffice.org employs a powerful BASIC-like programming language, which you can use to create your own functions. Although this language is called BASIC, it is several generations beyond the BASIC you might have used in the past. OpenOffice.org's BASIC is a high-level, object-oriented environment designed to appeal to programmers who wish to quickly add their own functions to the suite.

However, it's possible for any user to record a series of actions as a macro, which is then automatically turned into a simple BASIC program. This can be very useful if you wish to automate a simple, repetitive task, such as the insertion of a paragraph of text, or even something more complicated, such as searching and replacing text within a document.

To record a macro, select **Tools** ► **Macro** ► **Record Macro**. After you've selected this option, any subsequent actions will be recorded. All keyboard strokes and clicks of the mouse will be captured and turned automatically into BASIC commands. To stop the recording, simply click the button on the floating toolbar. After this, you'll be invited to give the macro a name (look to the top left of the dialog box). Once you've done so, click **Save**. You can then run your macro in the future by choosing **Tools** ► **Macro** ► **Run Macro**. Simply expand the **My Macros** and **Standard** entries at the top left of the dialog box, select your macro in the list, and then click **Run**.

Saving Files

As mentioned in Chapter 21, OpenOffice.org uses the OpenDocument range of file formats. The files end with an .ods, .odt, .odp, or .odb file extension, depending on whether they've been saved by Calc, Writer, Impress, or Base, respectively. The OpenDocument format is the best choice when you're saving documents that you are likely to further edit within OpenOffice.org. However, if you wish to share files with colleagues who aren't running Ubuntu or OpenOffice.org, the solution is to save the files as Microsoft Office files. To save in this format, just choose it from the Save As drop-down list in the Save As dialog box. If your colleague is running an older version of OpenOffice.org or StarOffice, you can also save in those file formats.

Alternatively, you might wish to save the file in one of the other file formats offered in the Save As drop-down list. However, saving files in an alternative format might result in the loss of some document components or formatting. For example, saving a Writer document as a simple text file (.txt) will lead to the loss of all of the formatting, as well as any of the original file's embedded objects, such as pictures.

To avoid losing document components or formatting, you might choose to output your OpenOffice.org files as Portable Document Format (PDF) files, which can be read by the Adobe Acrobat viewer. The benefit of this approach is that a complete facsimile of your document will be made available, with all the necessary fonts and on-screen elements included within the PDF file. The drawback is that PDF files cannot be loaded into OpenOffice.org for further editing, so you should always save an additional copy of the file in the native OpenOffice.org format. To save any file as a PDF throughout the suite, select File ► Export As PDF. Then choose PDF in the File Format drop-down box, as shown in Figure 22-5.




Figure 22-5. All the programs in the suite can export files in Adobe PDF format.

Summary

In this chapter, we looked at the configuration options provided with OpenOffice.org. You were introduced to the user interface, which is shared across all the programs within the suite, and learned how it can be customized. We also examined some common tools provided across the suite of programs, such as macro generation.

Over the following chapters, we will look at each major component of the suite, starting with Writer.



In Depth: Writer

The word processor is arguably the most popular element within any office suite. That said, you'll be happy to know that OpenOffice.org's Writer component doesn't skimp on features. It offers full text-editing and formatting functionality, along with powerful higher-level features such as mail merge.

In this chapter, we'll take a look at some of Writer's most useful features. As with all of the components in the OpenOffice.org suite, describing the features within Writer could easily fill an entire book. You should do some exploring on your own by clicking around to discover new features, as well as make judicious use of the help system.

Formatting Text

You can format text within Writer using several methods. Here, we'll look at using the Formatting toolbar, the context menu, and the Style and Formatting palette.

The Formatting Toolbar

Formatting text is easy to do via the Formatting toolbar, which is just above the ruler and main document area. Using the toolbar buttons, you can select the type of font you wish to use, its point size, and its style (normal, bold, italics, and so on). The range of fonts is previewed in the Font drop-down list, making it easy to select the right typeface.

In addition, the Formatting toolbar lets you justify text so that it's aligned to the left or right margin, centered, or fully justified. You can also indent text using the relevant icons. As with elsewhere in Ubuntu, a tooltip will appear over each icon when you hover the mouse cursor over it, as shown in Figure 23-1. To the right of the indentation buttons are tools to change the text background and foreground colors, and also a tool to create highlighter pen-style effects.




Figure 23-1. When you hover your mouse over an icon, a tooltip appears to explain what it does.

Context Menu

Rather than use the Formatting toolbar, you can format text using the context menu. Right-click the text you want to format, and a context menu will present options for the font, size, style, alignment, and line spacing. The context menu also allows you to change the case of the highlighted characters—from uppercase to lowercase, and vice versa.

By selecting the Character option from the context menu, you get ultimate control over the font formatting. This will present a dialog box that includes every possible option, such as rotating the text and altering the individual character spacing.

Tip The Character dialog box lets you create interesting typographical effects. The Paragraph dialog box has many options for formatting paragraphs. These tools open up the possibility of using Writer for simple desktop publishing work.

Selecting Paragraph from the context menu displays the Paragraph dialog box, as shown in Figure 23-2. This gives you control over paragraph elements, such as line spacing, indentation, and automatic numbering. Here, you will also find an option to automatically create drop caps, so you can start a piece of writing in style!

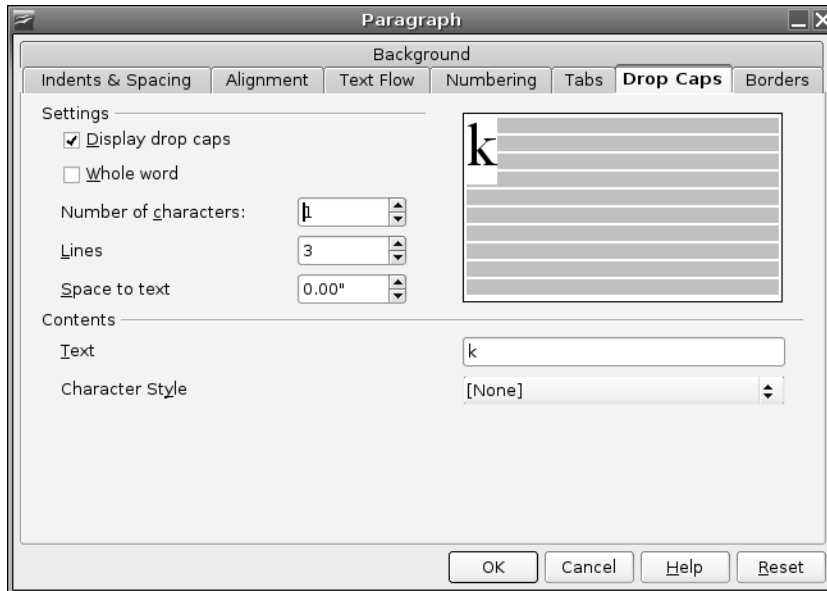


Figure 23-2. Writer includes many elements found in desktop publishing packages, such as the ability to create drop caps.

The Style and Formatting Palette

The Style and Formatting palette offers a variety of predefined formatting styles that you can apply to selected text or enable before you begin adding text. To make it appear, click **Formatting ► Styles and Formatting**, or press F11. You can simply click the palette's close button to get rid of it.

You can easily add your own text styles to the Style and Formatting palette. Simply select some text that has the formatting applied, click the top-right button (denoted by a paragraph symbol next to a block of text), and then select **New Style** from **Selection** in the list. You'll be invited to give the style a name, and when you click **OK**, it will appear in the list.

Spell-Checking

Writer is able to automatically spell-check as you type. Any words it considers misspelled will be underlined in red. You can choose from a list of possible corrections by right-clicking the word and selecting from the context menu. If you're sure the word is spelled correctly but it doesn't appear to be in the dictionary, you can select **Add ► Standard.dic** from the context menu, as shown in Figure 23-3. This will add the word to your own personal dictionary extension (other users won't have access to your dictionary and will need to create their own list of approved words).

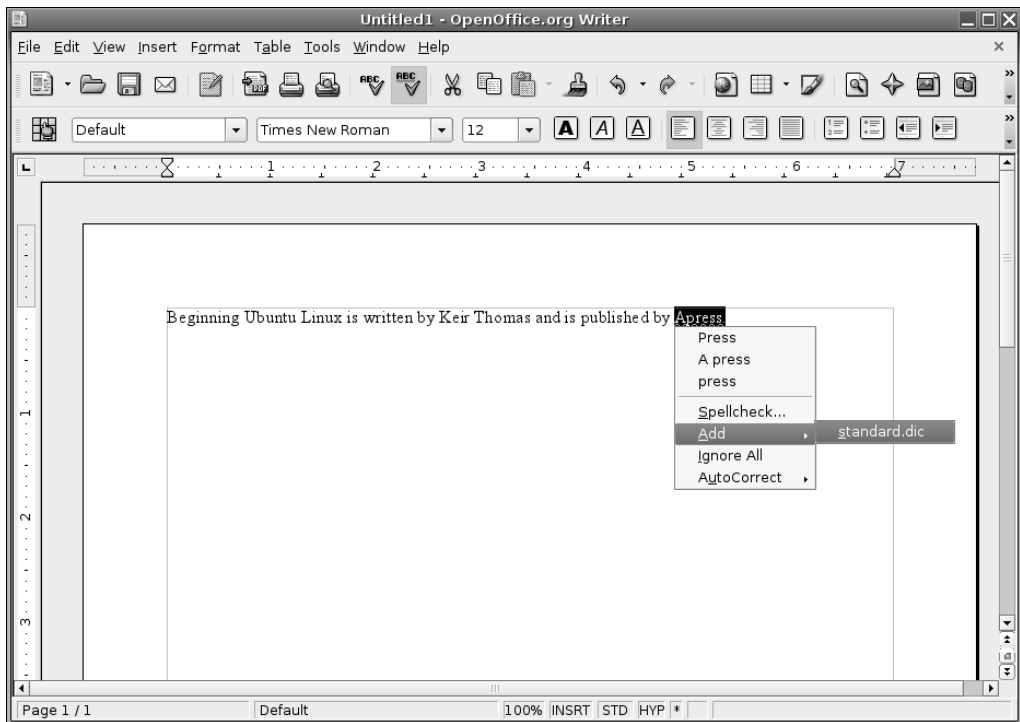


Figure 23-3. Any words you're going to use frequently, but which Writer doesn't recognize, can be added to your personal dictionary.

Tip You might find that the spell-checker is set for US English. If you live outside the United States, or need to create documents for readers in other countries, you can choose a dictionary tailored to your locality or needs. To change the language, select Tools ► Options. In the list on the left, select Language Settings, and then Languages. In the Default Languages for Documents list, select your local variation. This will then become the default for all new documents.

If you find live spell-checking invasive or distracting, you can deactivate it by selecting Tools ► Spellcheck, clicking the Options button, and removing the check next to Check Spelling As You Type.

You can manually spell-check the document at any time by clicking Tools ► Spellcheck. This will scan through the document and prompt you for corrections for words the program considers misspelled.

Inserting Pictures

Writer includes quite substantial desktop publishing-like functions, such as the ability to insert pictures into text documents and to have text flow around pictures.

Inserting any kind of graphic—a graph, digital camera photo, drawing, or any other type of image—is easy. Simply choose Insert ► Picture ► From File.

Tip If you have a scanner, you can also scan pictures directly into Writer documents. Simply click Insert ► Picture ► Scan ► Select Source.

After you've inserted a picture, you can place it anywhere on the page. When you select the picture, a new toolbar appears. This toolbar contains various simple image-tweaking tools, such as those for altering the brightness, contrast, and color balance of the image. Additionally, by clicking and dragging the blue handles surrounding the image, you can resize it.

Graphics that are imported into Writer must be anchored in some way. In other words, they must be linked to a page element so that they don't move unexpectedly. By default, they're anchored to the nearest paragraph, which means that if that paragraph moves, the graphic will move, too. Alternatively, by right-clicking the graphic, you can choose to anchor it to the page, paragraph, or character it is on or next to, as shown in Figure 23-4. Selecting to anchor it to the page will fix it firmly in place, regardless of what happens to the contents of the surrounding text. The As Character option is slightly different from the To Character option. When you choose As Character, the image will be anchored to the character it is next to, and it is actually inserted in the same line as that character, as if it were a character itself. If the image is bigger than the line it is anchored in, the line height will automatically change to accommodate it.

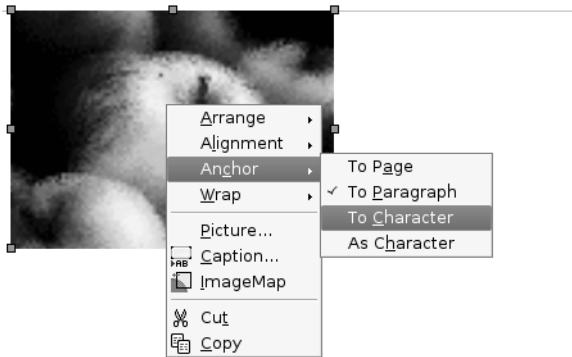


Figure 23-4. A picture can be “anchored” to the page, paragraph, or a character. This affects how it responds to the paragraphs surrounding it.

As you can see in Figure 23-4, the context menu also includes a Wrap option, which lets you set the type of text wrap you want to use. By default, Optimal Page Wrap is selected. This causes the text to wrap down just one side of the picture—the side on which the picture is farthest from the edge of the page. Alternatives include No Wrap, which will mean that the graphic will occupy the entire space on the page; no text is allowed on either side of it. However, Page Wrap is the best option if you’re looking for a desktop publishing-style effect, because the text will wrap around both sides of the picture. Alternatively, if you wish the image to appear in the background of the page with text flowing across it, you can select the relevant option from the context menu.

As always within OpenOffice.org, ultimate control is achieved by opening the relevant dialog box. You can set up how graphics are treated on the page by right-clicking the image and selecting Picture. In the dialog box that appears, you can select the wrap effect, specify the invisible border around the wrap (which governs how close the text is to the image), and give the image a border frame.

Working with Tables

Often, it’s useful to present columns of numbers or text within a word processor document. To make it easy to align the columns, OpenOffice.org offers the Table tool. This lets you quickly and easily create a grid in which to enter numbers or other information. You can even turn tables into simple spreadsheets, and tally rows or columns via simple formulas.

To insert a table, click and hold the Table icon on the Standard toolbar (which runs across the top of the screen beneath the menu). Then simply drag the mouse in the table diagram that appears until you have the desired number of rows and columns, and release the mouse button to create the table, as shown in Figure 23-5.

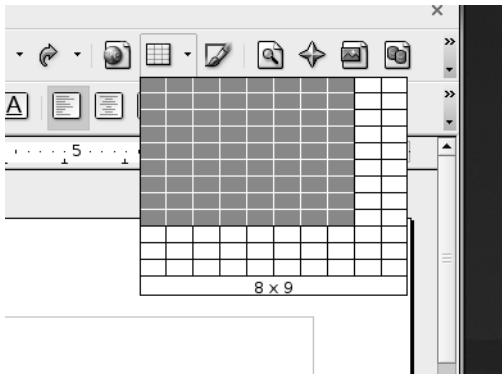


Figure 23-5. Just select the Table icon on the Main toolbar and drag the mouse to define the size of the table. Release the mouse button when you're finished.

Whenever your cursor is inside the table, a new toolbar will appear, offering handy options. Once again, simply hover your mouse over each button to find out what it does via a tooltip.

As with spreadsheets, tables consist of cells arranged into rows (running horizontally) and columns (running vertically). Altering the size of a column is easy. Just hover the mouse over the edge of a cell until it changes to a resizing cursor, and then click and drag. You can do the same on a horizontal bar to alter a cell's height, but a far better method is to right-click within the cell, select Row ► Height, and enter a value. This will ensure that subsequent cells are shifted down to make space for the newly enlarged cell, which doesn't happen when you click and drag the cell's border.

Tip An alternative way of resizing cells is to click in a cell and press the Enter key, which inserts a carriage return. Cells expand in size to fit their contents.

Once the cursor is within a table, you can move from cell to cell using the Tab key. Alternatively, you can move backward through the cells by pressing Shift+Tab.

To add more rows or columns, click the relevant icon on the Table toolbar (the fourth and fifth buttons on the bottom row). To split an existing cell, ensure your cursor is inside it, right-click, select Cell from the menu, and then click Split.

If you want to total figures within a table, click in an empty cell, and then click the Sum icon on the Table toolbar (the Greek sigma symbol on the right side of the bottom row). This is similar to inserting a function in a spreadsheet. The cell holds the formula for the sum and clicking additional cells, or a range of cells, adds them to the sum.

Note Only correctly formatted cells can be summed using the Sum icon on the Table toolbar. Cells with spaces or text within them cannot be added to the formula.

You can alter the styling of any cell using various icons on the Table toolbar, as well as the standard text formatting tools on the Formatting toolbar. The Table toolbar allows you to add borders to the cells and change the background colors. Alternatively, you can choose to remove all borders from the cells by clicking the Borders icon and then the No Borders option (note that gray borders will remain in place, but these are only for your convenience and won't appear in printouts).

Mail Merging

Mail merging refers to automatically applying a database of details, such as names and addresses, to a document, so that many personalized copies are produced. It's ordinarily used to create form letters for mailings.

OpenOffice.org makes the procedure very easy, but it requires source data that will be merged into the document. As with Microsoft Word, you can either enter this data within Writer itself or choose to import data from a separate document. Unless you have enough knowledge of databases to connect one to OpenOffice.org (the program works with dBase and MySQL files, among others), you may want to input existing data in the form of a comma-separated value (CSV) text file. This is the simplest form of data file that is understood by the majority of office programs and databases.

Here, we're going to look at entering the data within Writer, which is the best policy for smaller mail merge operations. You can then output the data as a CSV file, so you can use it again later. Here are the steps for using mail merge:

1. Select Tools ► Mail Merge Wizard to start the wizard, as shown in Figure 23-6.
2. Specify your starting document, which is the document in which the merged data will appear. You can opt to use the current document, create a new document, open a document from file, or use a template as the basis for your file.
3. Choose the mail merge type. You can choose to create a merged e-mail (for sending to multiple recipients) or a merged letter.
4. You're asked to tell Writer about your data. Writer needs to know where to find the addresses that will be merged into the document. Click the Select Address List button.
5. In the window that appears, you have a number of options. You can raid your Evolution e-mail address book for the data, click Add to select an already existing data source (such as a database or CSV file), or create a data source from scratch. Choose to create a data source to enter the data in Writer.

Note The fourth option for choosing a mail merge data source, Filter, allows you to filter the database source you select after clicking Add, so that you can import only specific data. To learn more about this technique, browse the OpenOffice.org Help file (click Help ► OpenOffice.org Help) and search for Filtering ► Data in Databases.

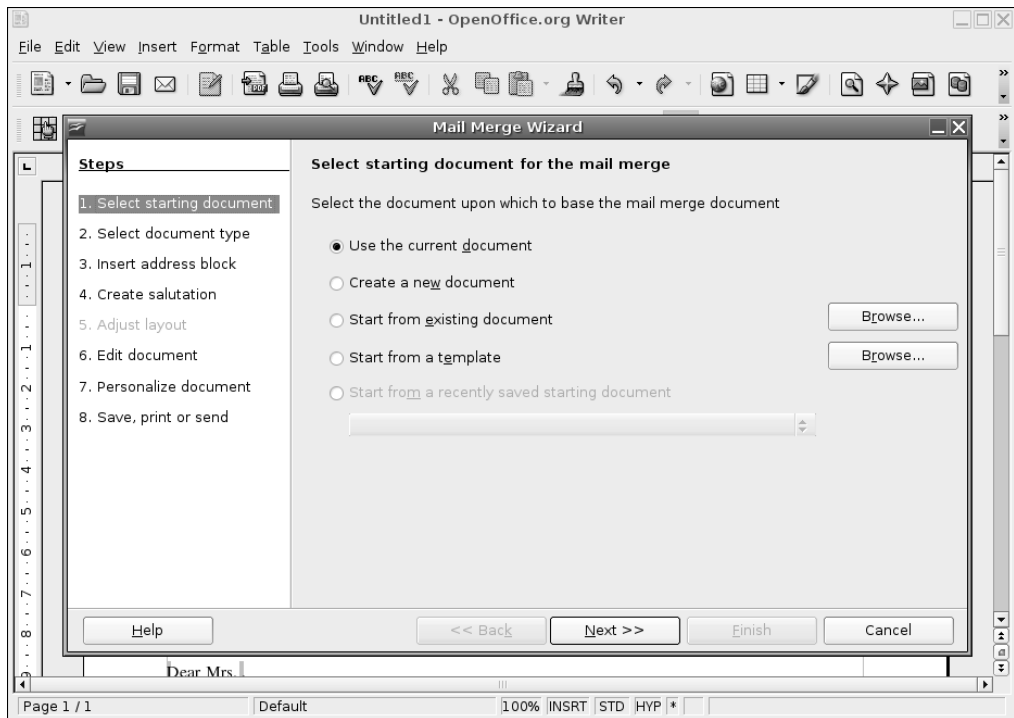


Figure 23-6. The Mail Merge Wizard makes creating multiple documents from a data source incredibly easy.

6. You're presented with a form for entering the data for each individual you want to receive the mail-merged letter, as shown in Figure 23-7. You don't need to fill in each field; you'll be able to choose which data fields to use in the document later on. If you wish to enter your own specific data types in addition to address details, you can click the Customize button to add your own field to the list. Using the up and down arrows in the window that appears, position the highlight where you would like the data to appear. Then click the Add button (alternatively, if there's a data field you're not using, you can highlight it and select Rename to reuse it). Obviously, you should add any new data fields you want *before* you begin to enter data!
7. Type in the data and press Enter at the end of each line. When you get to the last field, click the New button at the top right. When you've finished entering all the data, click OK. Then accept Writer's offer to save the data as a CSV file.
8. You're returned to the data-selection screen, and your just-saved file will be at the top of the list. Click OK.

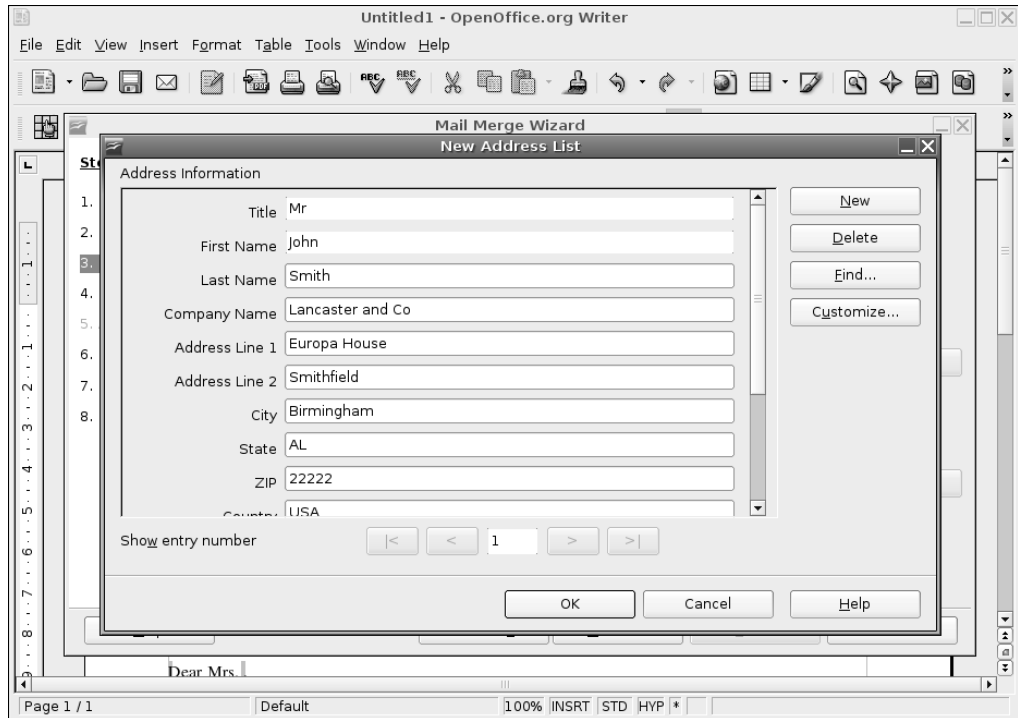


Figure 23-7. Enter the data for each person you want to receive the letter.

9. You're returned to the main Mail Merge Wizard window, where you can select whether or not to include an address block. All this means is that Writer will automatically add the merge fields to your document in what it considers the correct format (for example, title, followed by first and last name, with each line of the address underneath, and so on). You can insert the merge fields manually later on if you wish; in that case, remove the check from the "This document shall contain an address block" check box. Click Next.
10. You're invited to create the salutation that will head the letter. This will contain the merge data as well, so that you can personalize the letter. Again, you can accept the default, tweak it slightly, or choose not to have an automated salutation (so that you can create your own later).
11. Depending on your previous choices, and whether you accepted the automatic address block and salutation, you are now given the choice to adjust the layout of the document in a rough way or to actually edit it (note that even if you accepted the address block and/or salutation, you'll get a chance to edit the document in the next step anyway).

12. If you opt to edit the document, you can insert your choice of merge fields by clicking Insert ► Fields ► Other. Select the Database tab in the window that appears, and then select Mail Merge Fields on the left side of the window. Click the small plus symbol next to the data file you created earlier, which should be listed on the right, and you can then select and insert the merge fields. Once you've finished, click the Return to Mail Merge Wizard button.
13. Click Next to perform the merge. You're then given a chance to edit the actual mail-merged documents (which, depending on the quantity of data entries you created earlier, could number in the tens, hundreds or even thousands!).
14. You can save or print the *merged* document containing the data. To save the document creating the merge fields, click Save Starting Document.

Adding Headers and Footers

You may want to add headers and footers to long documents to aid navigation. They appear at the top and bottom of each page, respectively, and can include the document title, page number, and other information. Headers and footers are created and edited independently of the main document.

As you might expect, inserting both headers and footers takes just a couple of clicks. Select Insert ► Header or Insert ► Footer, depending on which you wish to insert (documents can have both, of course). Writer will then display an editing area where you can type text to appear in the header or footer. For more options, right-click in the area, select Page, and then click the Header or Footer tab. Here, you can control the formatting and nature of the header or footer. Clicking the More button will let you apply borders or background colors.

You might wish to insert page numbers that will be updated automatically as the document progresses. OpenOffice.org refers to data that automatically updates as a *field*. You can insert a wide variety of fields by selecting Insert ► Fields, as shown in Figure 23-8. For example, along with the page number, you can insert the document title and author name (which is read from the details entered into the Options configuration dialog box, accessed from the Tools menu). In addition, you can enter mail merge fields by clicking Other (see the previous section for a description of how to associate mail merge data with a document).

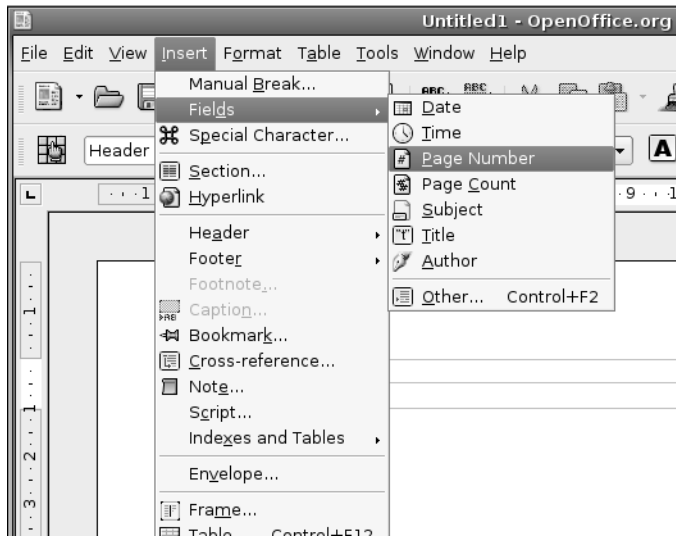


Figure 23-8. Automatically updating data, such as page numbers, can be inserted into headers and footers.

Summary

In this chapter, we've examined Writer, one of the core components of OpenOffice.org. We've looked at some of the key tools, which enable quick and easy document creation. In particular, you've learned how to format text, use the spell-checking component, insert pictures, create and edit tables, mail merge, and add headers and footers.

In the next chapter, we move on to another vital part of OpenOffice.org: Calc, the spreadsheet component.



In Depth: Calc

Calc is the spreadsheet component of OpenOffice.org. Like most modern spreadsheet programs, it contains hundreds of features, many of which few average users will ever use. However, it doesn't abandon its user-friendliness in the process and remains very simple for those who want to work on modest calculations, such as home finances or mortgage interest payments. In many regards, Calc is practically a clone of Excel, and anyone who has used Microsoft's spreadsheet program will be able to get started with it immediately.

In this chapter, you'll learn about some of the best features of Calc, as well as the basics of spreadsheet creation.

Entering and Formatting Data

As with all spreadsheets, entering data into a Calc document is simply a matter of selecting a cell and starting to type. You can enter practically anything into a cell, but a handful of symbols are not allowed. For example, you cannot enter an equal sign (=) followed by a number and have it appear in a cell, because Calc will assume that this is part of a formula.

Tip To enter any character into a cell, including an equal sign followed by a digit, precede it with an apostrophe ('). The apostrophe itself won't be visible within the spreadsheet, and whatever you type won't be interpreted in any special way; it will be seen as plain text.

Entering a sequence of data across a range of cells can be automated. Start typing the sequence of numbers, highlight them, and then click and drag the small handle to the bottom right of the last cell. This will continue the sequence. You'll see a tooltip window, indicating what the content of each cell will be. Figure 24-1 illustrates this process.

Cells can be formatted in a variety of ways. For trivial formatting changes, such as selecting a different font or changing the number format, you can use the Formatting toolbar. For example, to turn the cell into one that displays currency, click the Currency icon (remember that hovering the mouse cursor over each icon will reveal a tooltip). You can also increase or decrease the number of visible decimal places by clicking the relevant Formatting toolbar icon.

For more formatting options, right-click the individual cell and select Format Cells from the menu. This displays the Format Cell dialog box, where you can change the style of the typeface, rotate text, place text at various angles, and so on. The Border tab of the Format Cell dialog box includes options for cell gridlines of varying thicknesses, which will appear when the document is eventually printed out.

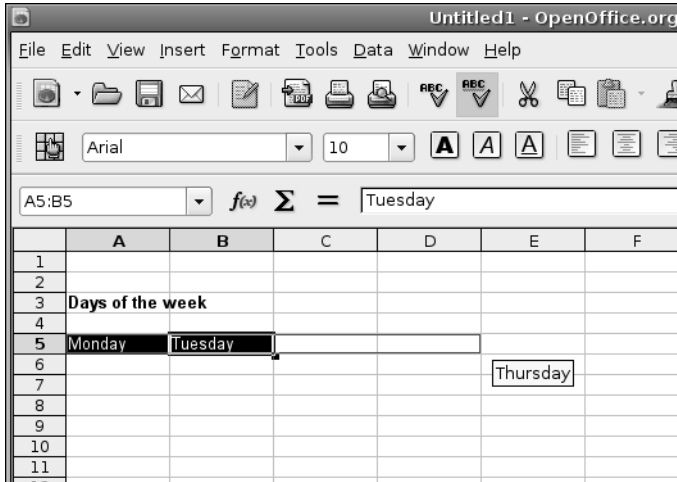


Figure 24-1. You can automate the entering of data sequences by clicking and dragging.

Deleting and Inserting Data and Cells

Deleting data is also easy. Just highlight the cell or cells with the data you want to delete, and then press the Delete key. If you want to totally eradicate the cell along with its contents, right-click it and select Delete. This will cause the data to the sides of the cell to move in. You'll be given a choice on where you want the cells to shift from to fill the space: left, right, above, or below.

To insert a new cell, right-click where you would like to it to appear and select Insert. Again, you'll be prompted about where you want to shift the surrounding cells in order to make space for the new cell.

Working with Formulas

Calc includes a large number of formulas. In addition to simple and complex math functions, Calc offers a range of logical functions, as well as statistical and database tools. Certain formulas can also be used to manipulate text strings, such as dates.

You can get an idea of the available functions by clicking the Function Wizard button on the Formula bar (which is just below the Formatting toolbar). This will bring up a categorized list of formulas, along with brief outlines of what function the formula performs. If you would like more details, use the help system, which contains comprehensive descriptions of most of the formulas, complete with examples of the correct syntax.

Note The Function Wizard is actually a continuation of the wizard system you've seen in other OpenOffice.org programs, and some of the functions are also available elsewhere in the suite.

You can reuse formulas simply by cutting and pasting them. Calc is intelligent enough to work out which cells the transplanted formula should refer to, but it's always a good idea to check to make sure the correct cells are referenced.

Using the Function Wizard

To use the Function Wizard to add a function, click the relevant button on the Formula bar, select the desired type of formula from the Category drop-down list, and then double-click an entry in the Function list to select it. Following this, you'll be prompted to input the relevant figures or define the appropriate data sources. Next to each text-entry box is a "shrink" button, which temporarily hides the wizard window, so you can select cells to be used within the formula.

Let's look at a quick example of using the wizard to work out an average value of a number of cells.

1. Select the cell in which you want the result of the formula to appear.
2. Start the Function Wizard by clicking the button on the Formula bar. In the left-hand list of functions, double-click AVERAGE. The wizard will then present a list of fields on the right side of the dialog box, where you can enter the values to be averaged. You could type numeric values directly into these fields, but it's more likely that you'll want to reference individual cells from the spreadsheet.
3. Click and drag the top of the dialog box to move it so that the spreadsheet underneath is at least partially visible.
4. Click the cursor in the first field of the dialog box, and then click the first cell you want to include in the calculation. This will automatically enter that cell reference into the field.
5. Click the next field in the dialog box, and then click the next cell you wish to include.
6. Repeat step 5 until all the fields you wish to include have been added to the fields in the dialog box (up to 30 can be selected; use the scroll bar on the right side of the wizard dialog box to reveal more fields).
7. Once you've finished, click the OK button. Calc will insert the formula into the cell you selected at the start, showing the result of the formula.

After you've added a formula with the wizard, you can edit it manually by clicking it and overtyping its contents in the Formula bar editing area. Alternatively, you can use the Function Wizard once again, by clicking the button on the Formula bar.

Summing Figures

To add the values of a number of cells, you could use the Function Wizard and select the SUM function, as shown in Figure 24-2. The procedure for choosing the cells is the same as described in the previous section.

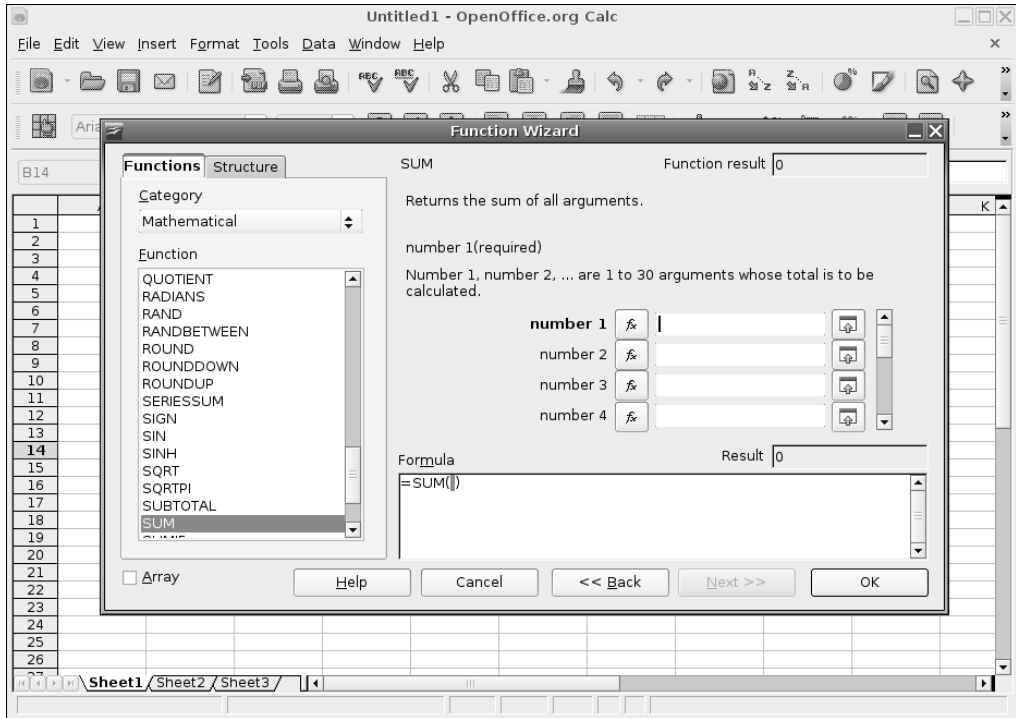


Figure 24-2. Creating formulas is easy using the Function Wizard.

However, Calc provides a far easier method of creating the sum formula. After positioning the cursor in an empty cell, simply click the Sum icon (the Greek sigma character) on the Formula bar, and then select the cells you wish to include in the sum. Then press Enter to see the results. If you place the cursor in a cell directly beneath a column of numbers, Calc may be clever enough to guess what you want to add and automatically select them. If it's incorrect, simply highlight the correct range of cells.

Tip You can select more than one cell by holding down the Ctrl key. You can select a range of cells in succession by clicking and dragging the mouse.

Sorting Data

Within a spreadsheet, you may want to sort data according to any number of criteria. For example, you might want to show a list of numbers from highest to lowest, or rearrange a list of names so that they're in alphabetical order. This is easy to do within Calc.

Start by highlighting the range of data you wish to sort. Alternatively, you can simply select one cell within it, because Calc is usually able to figure out the range of cells you want to use. Then select **Data** ► **Sort** from the main menu. Calc will automatically select a sort key, which will appear in the Sort By drop-down list, as shown in Figure 24-3. However, you can also choose your own sort key from the drop-down menu if you wish, and you can choose to further refine your selection by choosing up to two more sort subkeys from the other drop-down menus.




Figure 24-3. Data can be sorted so that it's in alphabetical or numerical order.

Creating Charts

Charts are useful because they present a quick visual summary of data. Calc produces charts through a step-by-step wizard, so it becomes very easy indeed. Here are the steps:

1. Highlight the data you want to graph. Be careful to include only the data itself and not any surrounding cells, or even the cell that contains the title for the array of data.
2. Select the **Insert** ► **Chart** menu option, or click the **Insert Chart** button on the Standard toolbar.

3. The cursor turns into a target with a small graph next to it. Click and drag on the spreadsheet itself to define the area of the graph. This can be any size. Also, you can resize it later.
4. The wizard starts. The first step is to define the range of cells to be used for the chart. By highlighting the cells before you started, you've already done this, so you can click the Next button. However, first make sure that the First Row As Label option is selected.
5. Choose the type of chart you wish to use. For most simple data selections, a bar graph is usually best. However, you might also choose to select a horizontal bar graph. Then click Next.
6. The wizard presents a subselection of graph types. You can also select whether gridlines are used to separate the various areas of the graph. Make your selections and click Next.
7. The last step allows you to give the chart a title and also choose whether you want a legend (a key that explains what the axes refer to) to appear next to it.
8. Click Create, and the chart will be created. Figure 24-4 shows an example.

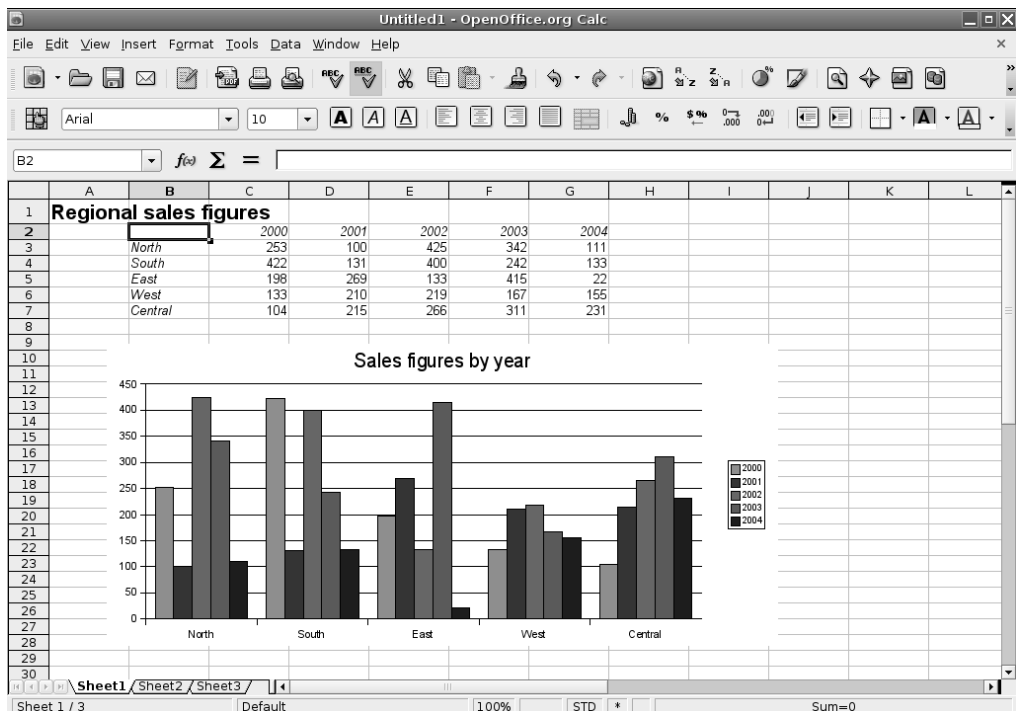


Figure 24-4. Creating a chart is easy within Calc and adds a professional flourish to your spreadsheet.

Once you've created a chart, you can alter its size by clicking and dragging the handles. You can also change various graphical aspects by double-clicking them. However, keep in mind that the graph is actually a picture, so the properties you edit are limited to changing the color and size of various elements.

The chart is linked to your data. Whenever your data changes, so will your chart. This is done automatically and doesn't require any user input.

Using Filters

The Filter function in Calc lets you selectively hide rows of data. The spreadsheet user then selects which of the rows of data to view from a drop-down list that appears in the cell at the top of the rows, as shown in the example in Figure 24-5.

	A	B	C	D	E
1	Regional sales figures				
2			2000	2001	2002
3		North	253	100	425
4		- all -	422	131	400
5		- Standard -	198	269	133
6		- Top 10 -	133	210	219
7		- empty -	104	215	266
8		- not empty -			
9		Central			
10		East			
11		South			
12		West			
13					
14					
15					
16					
17					
18					

Figure 24-5. Filters allow you to selectively hide or show rows of data in a spreadsheet.

Note A Calc filter is a little like an Excel pivot chart, especially when it's combined with an automatically generated chart.

Using filters in this way can be useful when you're dealing with a very large table of data. It helps isolate figures so you can compare them side by side in an easy-to-follow format. For example, you could filter a table of sales figures by year.

To use the Filter function, start by highlighting the data you wish to see in the drop-down list. Make sure the column header for the data is included, too. If you're using the Filter feature on a table of data, this selection can be any column within the table, although it obviously makes sense to use a column that is pertinent to the filtering that will take place. After you've selected the data to filter, select **Data** ► **Filter** ► **Autofilter**. You should find that, in place of the column header, a drop-down list appears. When a user selects a various entry in the list, Calc will display only the corresponding row of the spreadsheet beneath.

To remove a filter, select **Data** ► **Filter** ► **Hide Autofilter**.

Summary

In this chapter, we examined OpenOffice.org Calc. We looked at the basics of how data can be entered into a cell and how it can be formatted. Then you learned how to create formulas. This is easy to do with the Function Wizard function, which automates the task.

Next, you saw how to sort data in a spreadsheet. We also went through the steps for creating charts using a Calc wizard. Finally, we looked at creating data filters, which work rather like pivot charts in Microsoft Excel.

In the next chapter, we move on to Impress, the presentations component of OpenOffice.org.



In Depth: Impress

Impress is the presentation package within OpenOffice.org. At first glance, it appears to be the simplest of the key OpenOffice.org components, and also the one that borrows most the look and feel from Microsoft Office. However, delving into its feature set reveals more than a few surprises, including sophisticated animation effects and drawing tools. Impress can also export presentations as Macromedia Flash-compatible files, which means that many Internet-enabled desktop computers around the world will be able to view the files, even if they don't have Impress or PowerPoint installed on their computers.

In this chapter, you'll learn about the main features of Impress, as well as the basics of working with presentations.

Creating a Quick Presentation

As soon as Impress starts, it will offer to guide you through the creation of a presentation using a wizard. This makes designing your document a matter of following a few steps.

You'll initially be offered three choices: Empty Presentation, From Template, or Open Existing Presentation. When Impress refers to *templates*, it means presentations that are both predesigned and also contain sample content. Only two templates are supplied with Impress, so this option is somewhat redundant. However, you might choose to look at them later, if only to get an idea of what a presentation consists of and how it's made.

Tip When you become experienced in working with Impress, you can create your own templates or download some from the Internet. To create your own template, simply select to save your document as a template in the File Type drop-down list in the Save As dialog box. Make sure you place any templates you download or create in the `/usr/lib/openoffice2/share/template/en-US/present/` directory (you will need to have superuser powers to do this and should make sure the file permissions are readable for all users).

The standard way of getting started is to create an empty presentation. This sounds more daunting than it actually is, because the Presentation Wizard will start, asking you to choose from a couple of ready-made basic designs, as shown in Figure 25-1. You'll also be given a chance to choose which format you want the presentation to take: whether it's designed primarily to be viewed on-screen or printed out.




Figure 25-1. *The Impress Presentation Wizard guides you through the creation of a new presentation.*

After this, you'll be invited to choose the presentation effects, including the transition effect that will separate each slide when the presentation is viewed and the speed of the transition. If you wish, you can set the pause between slides, too, as well as the length of time each slide stays on the screen.

After clicking the Create button in the wizard, Impress will start, and you'll be invited to choose a layout for your initial slide. These are previewed on-screen on the right side of the program window. A variety of design templates are available, ranging from those that contain mostly text to those that feature pictures and/or graphs.

Depending on which template you choose, you should end up with a handful of text boxes on your screen. Editing the text in these is simply a matter of clicking within them. The formatting of the text will be set automatically.

Tip You can move and shrink each text box by clicking the handles surrounding the box. To draw a new text box, select the relevant tool on the Drawing toolbar, which runs along the bottom of the screen. Simply click and drag to draw a box of whatever size you want.

Working in Impress

When the Presentation Wizard has finished and Impress has started, you'll notice three main elements in the program window, from left to right, as shown in Figure 25-2. You work in these panes as follows:

Slides pane: This pane shows the slides in your presentation in order, one beneath the other. Simply click to select whichever slide you want to work on, or click and drag to reorder the slides. To create a new slide, right-click in a blank area on the Slides pane. Right-clicking any existing slide will present a range of options, including one to delete the slide.

Main work area: This is in the middle of the program window and lets you edit the various slides, as well as any other elements attached to the presentation, such as notes or handout documents. Simply click the relevant tab.

Tasks pane: Here, you can access the elements that will make up your presentation, such as slide templates, animations, and transition effects. Select the slide you wish to apply the elements to in the Slides pane, and then click the effect or template you wish to apply in the Tasks pane. In the case of animations or transitions, you can change various detailed settings relating to the selected element.

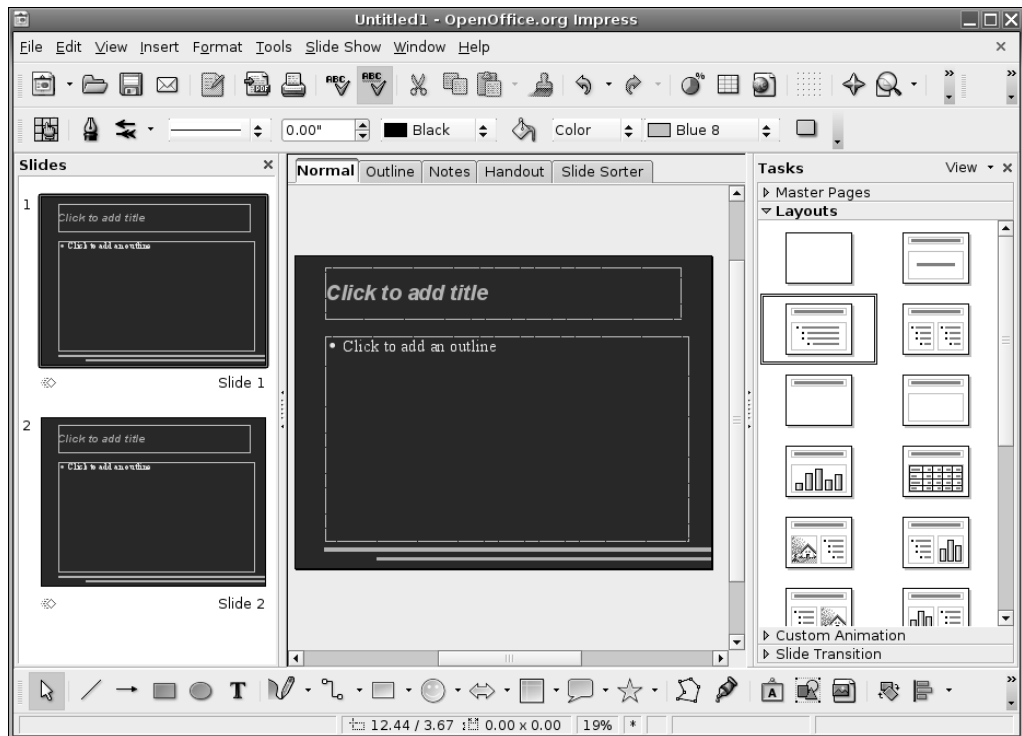


Figure 25-2. The main Impress window is split into three elements: the Slides pane, main work area, and the Tasks pane.

In addition, Impress has a Drawing toolbar, which appears at the bottom of the screen. This lets you draw various items on screen, such as lines, circles, and rectangles, and also contains a handful of special-effect tools, which I'll discuss later in this chapter, in the "Applying Fontwork" and "Using 3D Effects" sections.

You can hide each on-screen item by clicking the View menu and then removing the check next to it. Alternatively, by clicking the vertical borders between each pane, you can resize the pane and make it either more or less prominent on screen. This is handy if you wish to temporarily gain more work space but don't want to lose sight of the previews in the Slides pane, for example.

Animating Slides

All elements within Impress can be animated in a variety of ways. For example, you might choose to have the contents of a particular text box fly in from the edge of the screen during the presentation. This can help add variety to your presentation, and perhaps even wake up your audience!

Setting an animation effect is simply a matter of clicking the border of the object you wish to animate in the main editing area so that it is selected, selecting Custom Animation in the Tasks pane, and then clicking the Add button. In the dialog box that appears, select how you want the effect to work. As shown in Figure 25-3, you have four choices, each with its own tab within the dialog box:

Entrance: This lets you animate an appearance effect for the selected object. For example, you can choose to have a text box dissolve into view or fly in from the side of the screen. When you select any effect, it will be previewed within the main editing area.

Emphasis: This gives you control over what, if anything, happens to the object while it's on screen. As the name suggests, you can use this animation to emphasize various elements while you're giving the presentation. Some emphasis effects are more dramatic than others and this lets you control the impact. If you want to make an important point, you can use a dramatic effect, while more moderate information is presented with a more subdued effect.

Exit: As you might expect, this lets you add an exit animation to the object. You might choose to have it fly off the side of the screen or spin away off the top of the screen. The animation choices here are identical to the Entrance choices.

Motion path: This makes the selected element fly around on screen according to a particular path. For example, selecting Heart will cause the element to fly around describing the shape of a heart, eventually returning to its origin. A motion path is effectively another way of emphasizing a particular object.

Note You can apply only one effect at a time to an object, although several separate effects can be applied to any object.

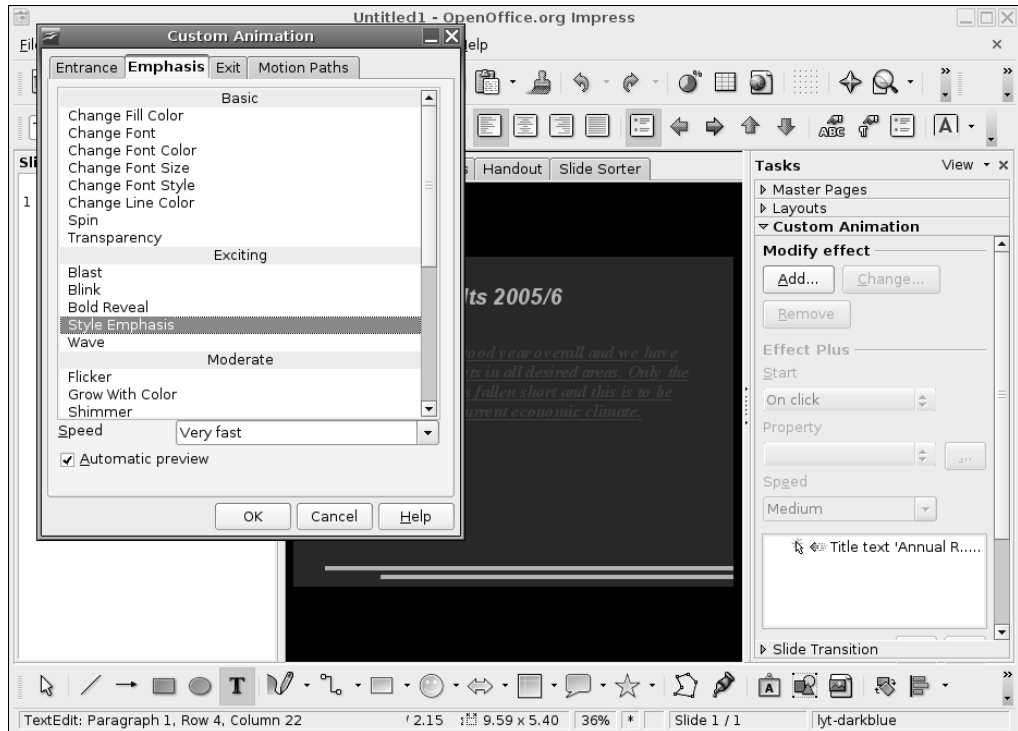


Figure 25-3. A wide variety of animation effects is available for on-screen elements.

With each animation, you can select the speed you wish it to play at, ranging from Very Slow to Fast. Simply make the selection at the bottom of the dialog box.

Once the animation has been defined and you've clicked OK, it will appear in a list at the bottom of the Custom Animation dialog box. You can choose to add more than one animation to an object by clicking the Add button again (ensuring the object is still selected in the main editing area). The animations will play in the order they're listed. You can click the Change Order up and down arrows to alter the order.

To fine-tune an effect, double-click it in the list to open its Effect Options dialog box (you can even add sound effects here). Under the Timing tab, you can control what cues the effect, such as a click of a mouse, or whether it will appear in sequence with other effects before or after in the list.

Applying Fontwork

The Fontwork tool lets you manipulate text in various playful ways, such as making it follow specific curved paths. You can find this tool on the Drawing toolbar, located at the bottom of the program window.

When you click the Fontwork icon, the Fontwork Gallery dialog box appears, offering a choice of predefined font effects. Don't worry if they're not quite what you want, because after you make a choice, you'll be invited to fine-tune it.

Once you've made the selection, the dummy text "Fontwork" will appear on screen. Editing the text is simple: just double-click the "Fontwork" text and type your own words. When you've finished, click outside the Fontwork selection.

Whenever the new Fontwork item is selected, a floating toolbar will appear, as shown in Figure 25-4. You can use this toolbar to alter various options. For example, you can select a completely different Fontwork selection from the gallery or, by clicking the second icon on the left, select your own path that you want the Fontwork item to follow.

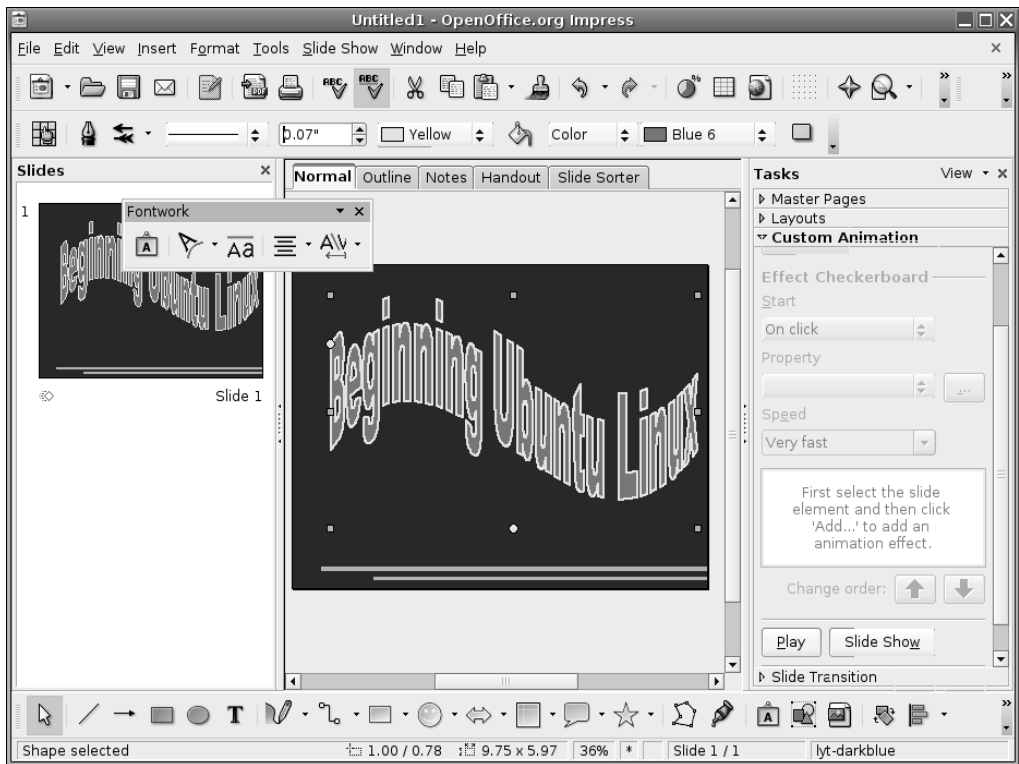


Figure 25-4. The Fontwork tool can add some special effects to your presentations.

You'll also see that the Formatting toolbar changes to allow you to alter the formatting of the Fontwork element. You can alter the thickness of the letter outlines, for example, or the color of the letters. Once again, the best way to learn how the tool works is to play around with the options and see what you can achieve.

To remove a Fontwork item, just select its border and press the Delete key on your keyboard.

Using 3D Effects

In addition to Fontwork effects, Impress includes a powerful 3D tool, which can give just about any on-screen element a 3D flourish (this tool is also available in some other OpenOffice.org applications). To use it, create a text box or shape using the Drawing toolbar at the bottom of the screen. Then right-click the text box or shape and select Convert ► 3D Object.

Note The 3D Object option is designed simply to give your object depth. If you want to create a genuine 3D object that you can rotate in 3D space, select the 3D Rotation Object.

You can gain much more control over the 3D effect by right-clicking it and selecting 3D Effects. This will open a floating palette window with five configuration panels, as shown in Figure 25-5. Click the icons at the top of the palette to adjust the type of 3D effect and its lighting, as follows:

Geometry: This defines how the 3D effect will look when it's applied to on-screen selections. For example, you can increase or decrease the rounded-edges value, and this will make any sharp objects on the screen appear softer when the 3D effect is applied.

Shading: This affects not the actual texture of the 3D object, but instead alters its color gradient. This is best demonstrated in action, so select the various shading modes from the drop-down list to see the effect. In addition, you can choose whether a shadow is applied to the effect, as well as the position of the virtual camera (the position of the hypothetical viewer looking at the 3D object).

Illumination: This lets you set the lighting effect. All 3D graphics usually need a light source because this helps illustrate the 3D effect; without a light source, the object will appear flat. Various predefined light sources are available. You can click and drag the light source in the preview window.

Textures: This affects how the textures will be applied to the 3D object. A texture is effectively a picture that is “wrapped around” the 3D object. Clever use of textures can add realism to a 3D object. A map of the world applied to a sphere can make it look like a globe, for example, or you could add wood or brickwork textures to make objects appear as tabletops or walls.

Material: This lets you apply various color overlays on the texture. This can radically alter the texture's look and feel, so it is quite a powerful option. To change the texture itself, right-click the object and select Area. This will present a list of predefined textures. Alternatively, you can select to use a color or pattern.

To apply any changes you make, click the check button at the top right of the palette. As with the other presentation effects, the best policy is simply to experiment until you're happy with the results.

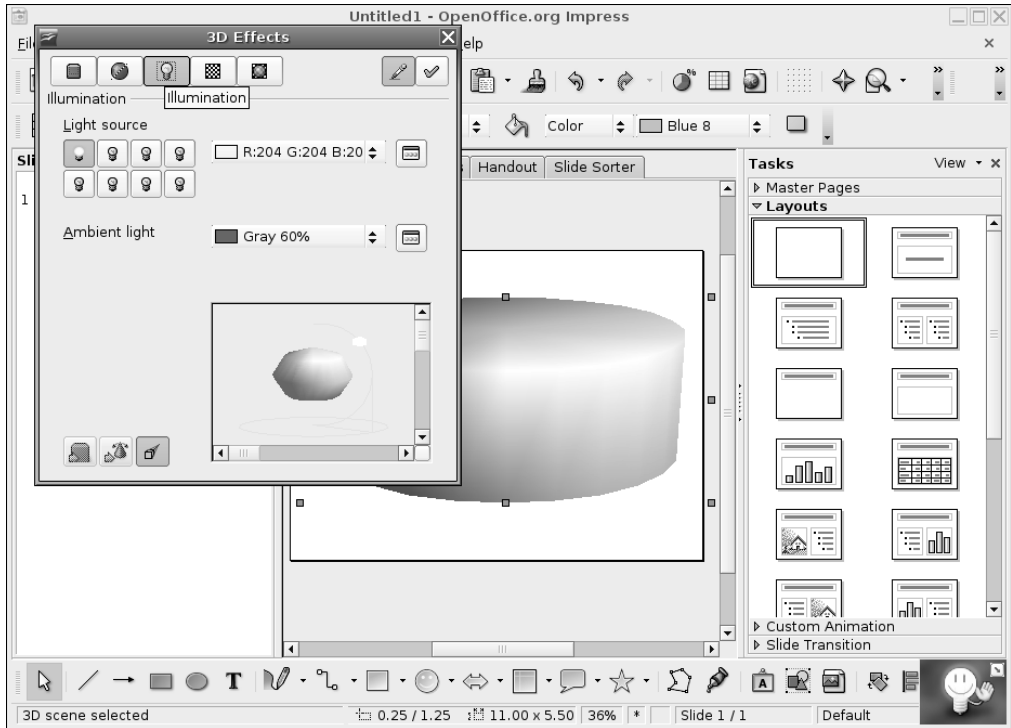


Figure 25-5. You can fine-tune 3D objects to quite a high degree using the 3D Effects palette.

Exporting a Presentation As a Flash File

If you plan to put your presentation online, or you want to send it to a colleague who doesn't have Impress or PowerPoint installed, outputting your presentation as a Flash animation could be a good idea. The process is simple. Just select **File** ► **Export**, and then select **Macromedia Flash (SWF)** in the **File Format** drop-down list (SWF is the Flash file type, which stands for Shockwave Flash). No further configuration is necessary.

In order to play the file, it needs to be opened within a web browser that has the Flash Player installed. This can be done by selecting **File** ► **Open** on most browsers, although you can also drag-and-drop the SWF file onto the browser window under Microsoft Windows. There shouldn't be much of a problem with compatibility, since the Flash Player is ubiquitous these days. If the web browser doesn't already have Flash installed, it's easy to download and install it (see www.macromedia.com/go/getflashplayer).

When the Flash file is opened in a web browser, the presentation starts, as shown in Figure 25-6. You can progress through it by clicking anywhere on the screen.

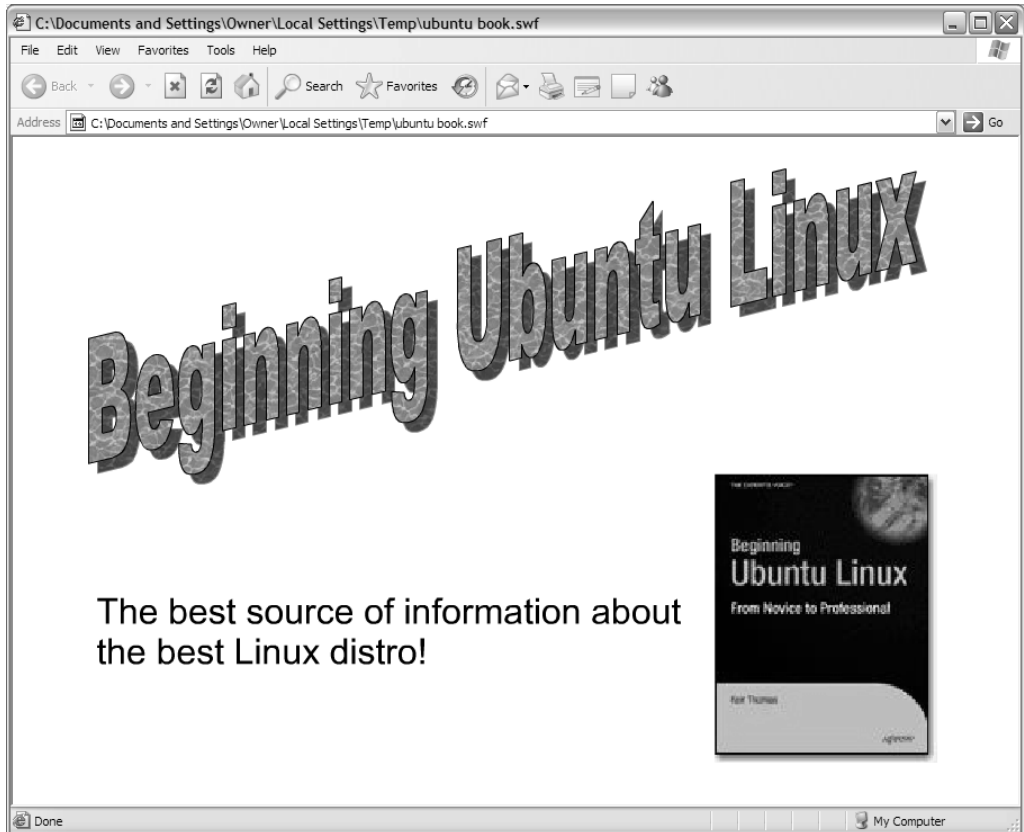


Figure 25-6. You can save any presentation as a Flash animation, which can be played back in a suitably equipped web browser.

Summary

In this chapter, we examined Impress, which is the presentations component within OpenOffice.org. We started by looking at how you can use the Presentation Wizard function to automate production of a basic Impress document. Then you saw how various effects can be added to the presentation, including 3D effects. Finally, we looked at how the presentation can be exported as a Shockwave Flash file for playback on virtually any web browser.

In the next chapter, we will explore the database component within OpenOffice.org: Base.



In Depth: Base

OpenOffice.org includes a number of tools to both interface with database servers and perform tasks such as enter and edit data. However, for most day-to-day users who have humble needs, creating such a setup is rather complicated. It requires some knowledge of how databases work on a technical level. For this reason, a new component was added to OpenOffice.org 2.0: Base.

Base is a relational database along the lines of Microsoft Access and is perfect for database applications of all sizes, including more modest efforts. For example, you could use it to create an inventory database to produce a report showing all products added for a certain geographical region on a certain date, or you could use it to catalog items in your personal stamp collection.

Relational databases such as those created by Base are ideal for quickly creating catalogs of information, such as inventory lists. In addition to making database creation simple and quick, relational databases let you easily query data to produce reports tailored to individual needs.

Base works on a number of levels depending on the knowledge of the user, but in its most basic form, it offers a design-based approach to the creation of tables and forms. Anyone who has previously created a database under Access will feel right at home.

In this chapter, we'll work through an example of using Base to create a simple database cataloging a collection of music. You can use the same techniques to create any kind of relational database.

Getting Started with Base

You'll find Base under the Applications ► Office menu. When the program first starts, the Database Wizard guides you through either creating a new database or opening an existing one, as shown in Figure 26-1.

The first step in creating a new database is to register it within OpenOffice.org. This means that it will be made available in other OpenOffice.org programs, such as Calc or Writer. Although the knowledge needed to use a database in this way is quite advanced, there's no harm in agreeing to this option. It might prove useful in the future as you learn more about OpenOffice.org.

Following this, you can choose to open the database for editing and/or start the Table Wizard. Once you click the Finish button, you'll be invited to give the database a name and save it immediately.



Figure 26-1. Base starts with the Database Wizard to facilitate the quick and easy creation of new databases.

Note Databases aren't like other office files in that they automatically save and update themselves. When using the finished database, you can simply enter data and then quit the program, without needing to deliberately opt to save the file.

Assuming that you did not opt to go directly to the Table Wizard, you'll now see the main Base program window. The right of the Base program window contains the Tasks and Data panes. The Tasks pane allows you to create new database elements, depending on what's selected in the Database pane. The Data pane shows any elements you've already created. The Database pane contains the four elements of the database that you can create and manipulate:

Tables: A table is what holds the actual data you'll eventually input. Therefore, a table is the first thing that needs to be created within a new database. Creating a table involves defining the types of data that you want to store and giving the individual data fields easy-to-understand names.

Forms: Although it's possible to enter data directly into a table, it isn't particularly intuitive or easy. Because of this, forms are used to make the data fields presentable. As the name suggests, in terms of layout these are not unlike the paper forms that you fill in to facilitate the collection of data by businesses. As with tables, forms must be created from scratch in a new database. Forms have *controls*, which are used to facilitate data entry, or to allow users to navigate the database or otherwise manipulate it. The most common type of control is a text-entry field, which is then tied to a data field within the table, but you can also have controls that perform certain functions, such as deleting a record in the database.

Queries: A query is a way of filtering the database so that you see only a subset of it. For example, in a database detailing sales figures from across the country, you might create a query to show only the data from a particular state.

Reports: A report is a way of presenting data for human consumption, usually in a printed format. For example, you could create a report that details sales figures in the form of a letter, or you might make a report to produce address labels using addresses stored in the database.

Note The usefulness of both queries and reports are that they can be saved and used over and over again, so you could use the same query each month to examine just a small section of the data. Base offers wizards to automate the creation of both queries and reports.

Double-clicking an item in the Database pane displays or activates that item. Right-clicking a Database pane item displays a variety of options related to editing the file.

Now, let's work through an example of using Base. First, you'll create a table, and then you'll create a form.

Creating a Database

As an example of using Base, you'll build a database, ready for data entry. The first step in the creation of a database is to make a table. This will hold the data that you will eventually enter using a form.

Adding a Table

As with all components within Base, you can use a wizard to create the table. The Table Wizard offers a number of predefined data fields corresponding to typical databases. It is fine for general use, but if you have a specific and unusual database in mind, you will need to create the table manually.

Here, you'll create a database to catalog CDs. This is easily accomplished with the Table Wizard, as follows:

1. Click the Tables icon in the Database pane, and then click the Use Wizard to Create Table icon. The Table Wizard starts.
2. You're given a choice between creating a business or personal database. As you would expect, business databases are likely to contain fields relating to business matters, such as accounting, and the fields in the personal section relate more to domestic matters. Choose Personal for this example.
3. Choose an entry from the Sample Tables drop-down list. For this example, select CD Collection.

4. In the Available Fields box, you now see a number of data fields that would prove handy for a CD collection. You don't need to use all of these. Instead, select only those you want in your table, and then click the single right-facing arrow button to transfer them to the Selected Fields box. For this example, select AlbumTitle, Artist, ReleaseYear, and Review, as shown in Figure 26-2. Then click the Next button. (Don't worry if you find the fields lacking or if you want to add your own—you'll see how to do just that in step 6).

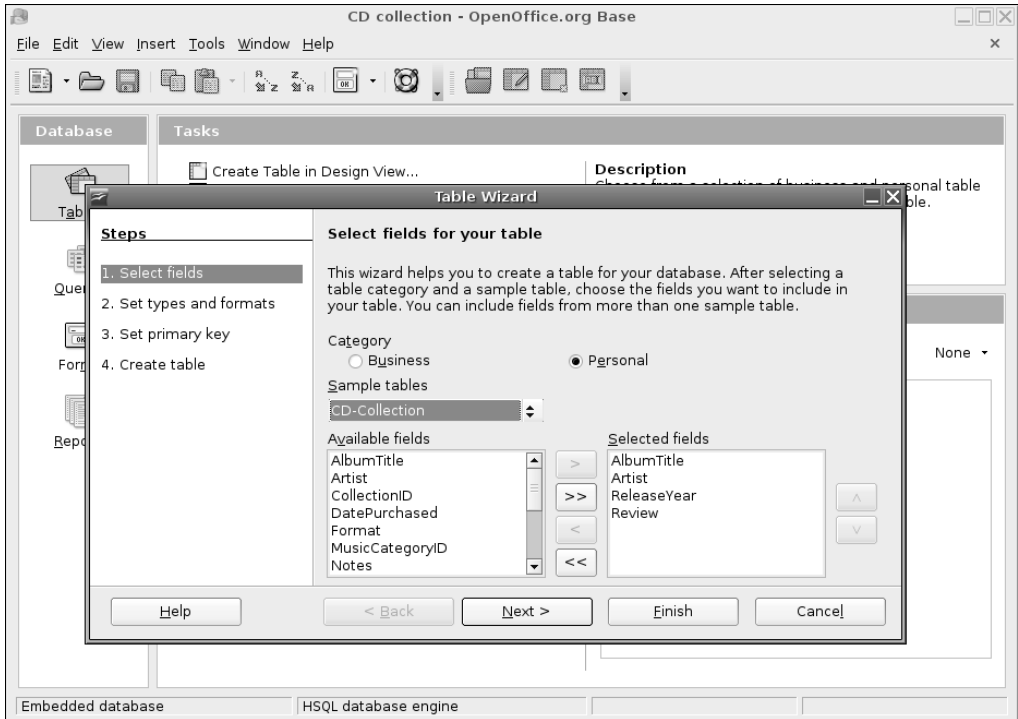


Figure 26-2. The Table Wizard contains ready-made data fields for a wide-variety of uses.

5. Check to make sure the fields you selected are of the correct type. Click each to see the information in the right area of the dialog box. Fields can take various forms depending on what kind of data they're supposed to hold. For example, one field might be designed to contain text, while another might need to contain numbers. Yet another might need to contain dates, and some can even contain pictures. As you might expect, the wizard has automatically selected the correct data types for the predefined fields.
6. For this example, you want to add a check box that shows whether the CD is scratched. If the CD in question is scratched, the user can click in a check box. If the CD isn't scratched, the box can be left blank. To create a check box, you need a special kind of data field called a Boolean. This means that the data field can be either true or false or, to put it a simpler way, it can hold either yes or no. To create a yes/no data field, click the plus button at the bottom of the Selected Fields box. This allows you to add another field. In the Field Name box, type **Scratched**. For the Field Type, click the drop-down list and locate the entry marked Yes/No [BOOLEAN]. The other options can remain as they are. Click Next to continue.

7. You're asked if you want to create a primary key. This is the unique numeric field that the database uses to keep track of each entry in the database. It's a must in a database like this one. The default choices are correct, so you can click Next again.
8. You've completed the Table Wizard. The next step is to create a form, so select Create a Form Based on This Table, and then click the Finish button.

Creating a Form

Forms are present in databases for the benefit of users to facilitate the quick-and-easy entry of data. They present data fields that you've just created within the table in an easy-to-understand form.

Base is able to walk you through the creation of forms via the Form Wizard. If you didn't select to run the Form Wizard previously, you can start it by clicking Forms in the Database pane, and then clicking Use Wizard to Create Form. Then follow these steps:

1. In the Form Wizard's first step, select which fields you want to appear on the form. As with the Table Wizard, this is simply a matter of selecting the fields, and then clicking the right-arrow button so that they appear under the Fields on the Form heading. Alternatively, by clicking the double-arrow button, you can select all of them in one fell swoop, which is what you want for this example.
2. You're asked if you want to create a subform. As its name suggests, this is effectively a form within your main form. A subform is useful with more complicated databases, where it might be necessary to view other data while filling in the form. For this simple example, leave the Add Subform box unchecked and click Next.
3. Choose a general layout for the data fields. The default is the table view, which many find ugly, so you might choose one of the first two options (in my database I chose the second option, as can be seen in Figures 26-3 and 26-4). These arrange the data fields in a spacious manner and make the form much more usable. If you look behind the wizard dialog box, you'll see a preview of how the form will look.
4. You're asked whether or not you want existing data to be displayed on the form. You can choose to treat the form as one created only for entering new data, so that you can't use it to navigate through the database and see existing data you've already entered. This might be useful in applications where you don't want users to see the other data in the database. However, for a database for your own personal use, being able to see the existing data is very handy, which is why The Form Is to Display All Data option is selected by default. For this example, simply click the Next button to accept the default.
5. Choose a look and feel for your form from the variety of color schemes available, as shown in Figure 26-3. Again, you can see them previewed behind the wizard dialog box. Feel free to experiment with the options under the Field Border heading. I prefer the 3D Look option, which gives the form elements a slight interior shadow, a common feature on most modern user interfaces. The Flat option simply adds a black border to the boxes, and the No Border option removes the border completely.

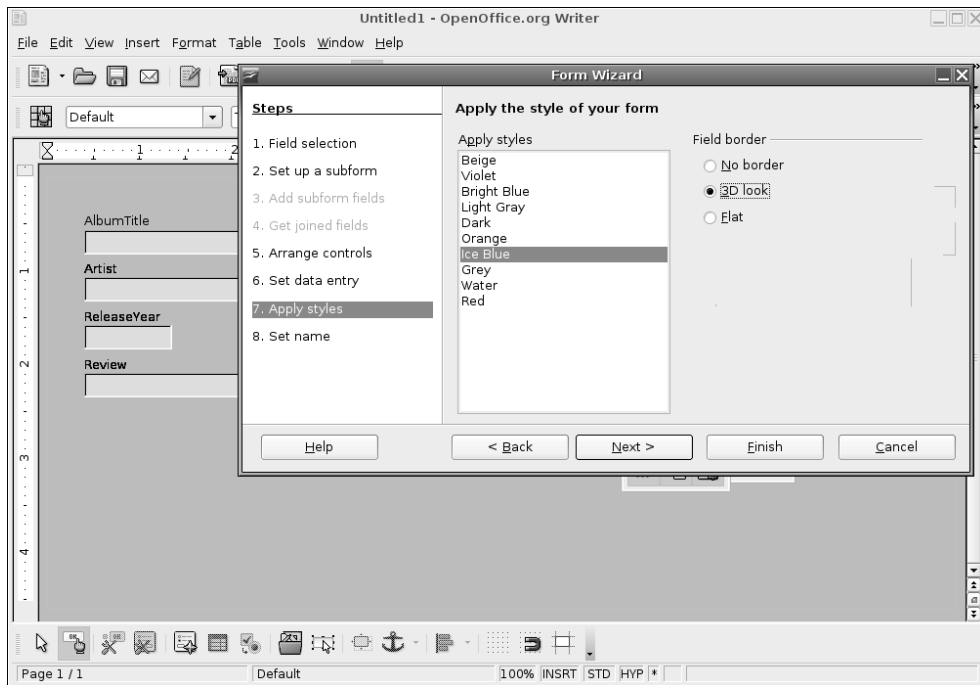


Figure 26-3. You can choose from a variety of look and feel options for your form, and each will be previewed behind the wizard dialog box.

6. You're invited to give the form a name. Enter a suitable name, such as **CD-Collection**. You are also given the option of entering data directly into the form or modifying it manually. You might have noticed that the field you created, the Scratched yes/no check box, isn't present on the form. Therefore, you need to add it to the form, so select to modify it manually. You'll do that in the next section.

Note There are no rules governing form names, and you can use virtually any symbols and also insert spaces into the name. However, it's a good idea to keep the form name simple and concise.

Adding Controls to the Form Manually

When the Form Wizard finishes, you should find yourself editing the form directly. A floating palette—the Form Control toolbar—will appear, offering various form-specific functions. Follow these steps to add the check box field to your form:

1. On the Form Control toolbar, click the icon that looks like a check box (on my computer, it was the third icon from the left), and then click and drag to draw a check box on the form. You need to make it big enough so that the label can be seen. If you release the mouse button too early, simply click and drag the handles at the edges to resize the control.

2. Once the check box has been drawn, double-click it. This will open Properties dialog box. Click in the Label box, delete what's there already, and type **Scratched?** (or anything else that will help you identify the check box on the completed form).
3. Click the Data tab and, in the Data Field box, select Scratched from the drop-down list, as shown in Figure 26-4.




Figure 26-4. Creating a custom control is simply a matter of drawing it on the form, and then matching it with a data field.

4. Close and save the form.

You can add more custom controls following the same basic approach you used here. Simply draw them on to the form, and then match them up with an entry in the table using the Data tab.

The database is almost ready for use. You just need to take one more step to modify the table, which you'll do next.

Editing the Database Table

Before you can use the database, you need to make a small change to the table you created earlier. Although the Table Wizard created a primary key, it didn't make it into an automatically updating number. Without this option activated, the user will need to manually number each entry in the database as it's created.

Follow these steps to edit the table and activate automatic numbering for the ID field:

1. Click the Tables icon in the Database pane of the main program window, right-click the table you created earlier, and select Edit.
2. Look for the entry in the table list labeled ID. It should be first in the list. Make sure that the cursor is on the ID line, and click the Autovalue drop-down list at the bottom of the window. Make sure that it reads Yes, as shown in Figure 26-5.




Figure 26-5. Set the primary key to automatically update by editing its value in the table.

3. Close the window and opt to save the table.

That's it! Your database is now ready to use.

Using the Database

Entering data into the finished database is easy. Click the Forms icon in the Database pane, and then double-click the form you created earlier.

The Form Control toolbar will still be visible. To hide it, click its close button at the top-right corner of the toolbar. After this, you can start to enter data into the form, as shown in the example in Figure 26-6. Note that you do not need to enter data in the ID field, because this will automatically be filled with the primary key number.




Figure 26-6. The finished database form lets you enter data into the input fields and navigate using the toolbar at the bottom.

Once you've filled in the form, you can click the Next Record button in the Form Navigation toolbar running along the bottom of the window (don't forget that hovering the mouse cursor over each button reveals a tooltip explaining what the button does). This will move you on to the next blank form, where you can enter more data. Repeat this as many times as necessary.

The Form Navigation toolbar contains other handy tools. For example, the first button—a magnifying glass—lets you search the database for a particular entry. It's well worth investigating the functions.

Summary

In this chapter, we looked at the Base database component of OpenOffice.org and how to use it to easily create and edit simple databases. We stepped through an example of setting up a database table and creating a database form that users can employ to enter and edit data.

In the next chapter, we will look at how to run Microsoft Office and many other Windows programs under Ubuntu.



In Depth: Evolution

Evolution isn't part of the OpenOffice.org suite. It was originally created by Ximian, an organization founded by the creators of the GNOME Desktop Project and acquired by Novell in August 2003. Even long after the acquisition, Evolution is still developed by many of the key GNOME desktop developers.

Although it's not explicitly described as such by its developers, Evolution is considered the "official" GNOME desktop e-mail program. Nearly every Linux distribution that uses the GNOME desktop system also uses Evolution. Evolution even retains the same look and feel as many elements of the Ubuntu desktop.

In terms of functionality, Evolution is similar to Microsoft Outlook, in that as well as being a powerful e-mail client, it incorporates contacts management, a calendar, and a to-do list. Evolution is even able to connect to Microsoft Exchange (2000 and above) groupware servers and synchronize with contact and calendar data, in addition to fetching e-mail. Of course, it can also connect to standard POP3/SMTP e-mail servers, as well as IMAP and also Novell GroupMail servers. This means it is compatible with practically every e-mail system in use today.

Although Evolution offers many of the functions of Microsoft Outlook, it differs in some key ways. Therefore, this chapter describes how to perform basic tasks, as well as more advanced everyday jobs.

Evolution Modes

Evolution consists of four components: Mail, Contacts, Calendars, and Tasks. These are interconnected but operate as separate modes within the program. Each mode can be selected using the switcher located at the bottom-left side of the program window. Simply click the button for the mode you wish to use. The program window, toolbar, and menu system will change to accommodate whichever mode is selected. Figure 27-1 shows the program in Mail mode.

Tip You can shrink the switcher component to small icons by clicking View ► Switcher Appearance ► Icons Only.

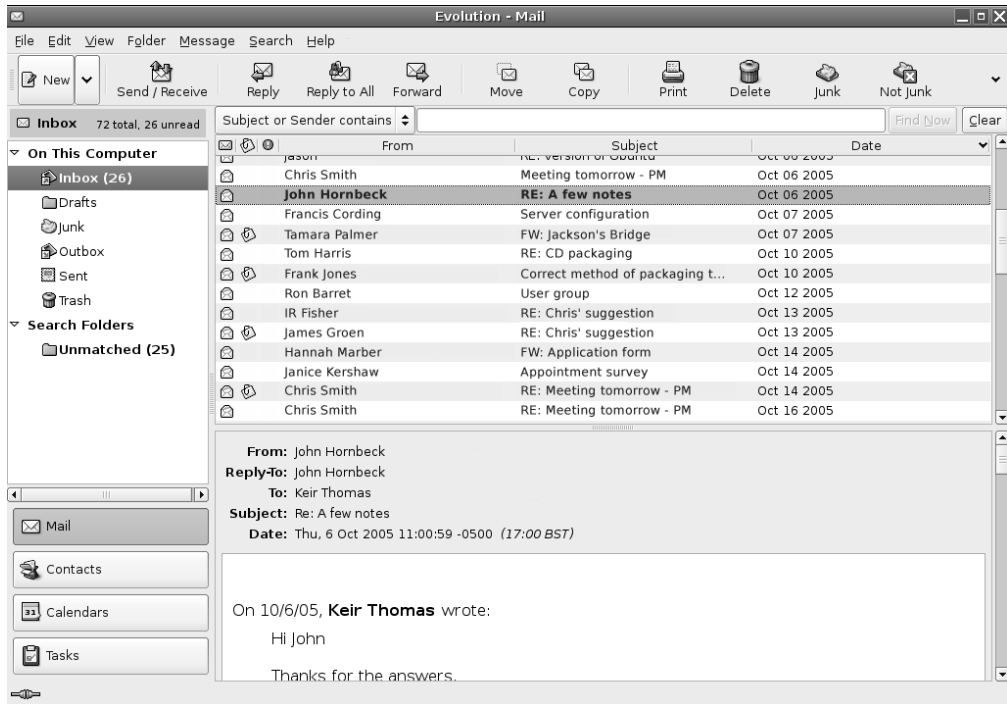


Figure 27-1. You can switch between Evolution's modes by clicking the buttons at the bottom left of the program window.

The four Evolution modes work as follows:

Mail: At the top left of the program window are the mail folders. Here, you'll find the Inbox and Sent folders, along with any other mail folders you create. On the right is the list of e-mail messages, and beneath this is the message preview pane, where the body of any message you select will be displayed. Above the message list is the search box, which works like most e-mail search routines: type the relevant word(s) and then click Find Now. Notable icons running along the top of the window include the New button, which will let you compose an e-mail message, and the Send/Receive button, which will download new messages and also send any messages in the Outbox folder.

Contacts: At the left side of the program window are listed the various contact folders. For most users, there will be just one, named Personal, but if you specified a groupware server during setup, you will also be able to connect to this by clicking its entry. At the top right is the list of contacts. Clicking any contact displays that individual's information at the bottom of the window, in the contact information area. The search bar at the top of the window beneath the toolbar lets you quickly search for contacts using their name. The New button on the toolbar lets you create a new contact.

Calendars: On the left side of the program window are the various calendars you can access. For most users, the Personal calendar will be the principal one, but you can also access shared calendars here. To the right is the monthly calendar and, in the middle of the program window, the appointment list, with half-hour entries covering the working day. By default, the current day is shown. To select a different day, simply double-click the day in the month view. You can switch between day, week, and month appointment views by clicking the Day, Work Week, and Month buttons on the toolbar.

Tasks: Your tasks lists are listed on the left side of the program window. Once again, most users will use just the Personal task list. The task list itself appears in the main program window.

Basic E-Mail Tasks

Evolution's e-mail functionality is arguably the heart of the program. Although it offers a lot of features, it is quite simple to use. If you've ever used any other e-mail client, such as Microsoft Outlook, you have a head start.

This section describes how to accomplish several everyday tasks within the e-mail component of Evolution. When you start Evolution, the e-mail mode is selected automatically. However, if it isn't, or if you've switched to a different mode within the program, simply click the Mail button at the bottom right of the program window.

Sending and Receiving E-Mail

Once Evolution has been set up correctly to work with your e-mail servers, as outlined in Chapter 8, you can simply click the Send/Receive button on the toolbar to connect to the server(s) and both send and receive e-mail.

You may need to enter your password if you didn't enter it during setup. You can check the Remember Password box to avoid having to type it again, but this will mean the password is then stored on your hard disk, possibly posing a security risk.

Note Although e-mail is normally sent as soon as you click the Send button when composing e-mail, if the sending has been delayed for any reason (such as because you were offline at the time), it will be sent as soon as you click the Send/Receive button. Until that point, it will be held in the Outbox folder on the left side of the program window.

Any outstanding mail is sent first, and then the receiving procedure is started. As shown in Figure 27-2, a status dialog box will tell you how many messages there are and the progress of the download. Clicking the Cancel button will stop the procedure (although some messages may already have been downloaded).

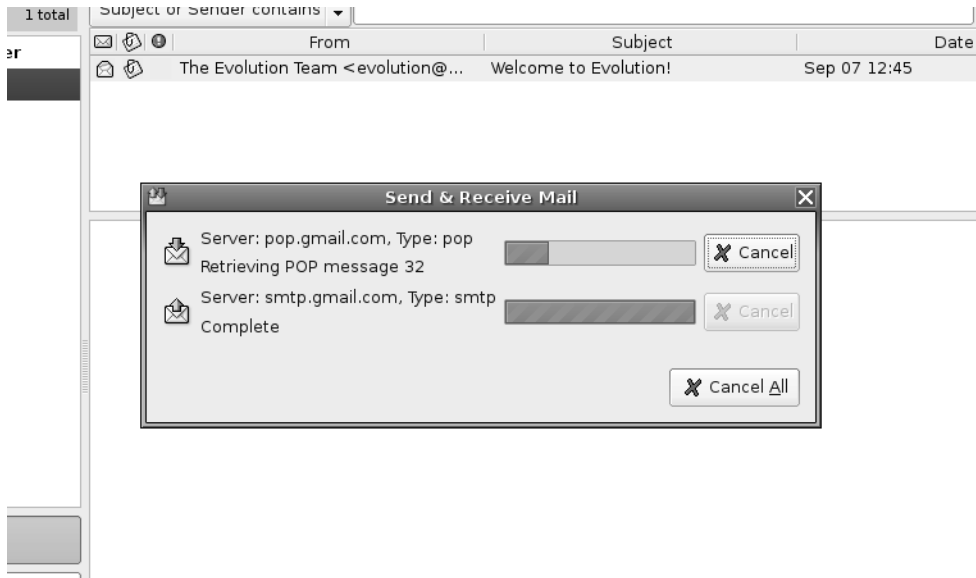


Figure 27-2. You'll see how many messages there are to download whenever you click the Send/Receive button.

Reading E-Mail

Simply click an e-mail message to view it in the preview pane at the bottom of the screen. Alternatively, you can double-click a message to open it in its own program window (selecting a message and pressing Enter will have the same effect).

As with most e-mail clients, any unread messages in the list appear in bold, and messages that have been read appear in ordinary type. By default, each message is marked as read after 1.5 seconds, but you can alter this value. To change it, click Edit ► Preferences, click the Mail Preferences icon in the Preferences dialog box, and then change the value under the Message Display heading. A value of 0 will cause the mail to switch to read status as soon as it's clicked, which can be useful if you want to quickly clear a lot of messages.

You can also mark many messages as read by highlighting them all, right-clicking an individual one, and selecting Mark As Read from the menu that appears. You can select multiple messages in the usual way: Shift-click to select a consecutive list or Ctrl-click for nonconsecutive selections.

Deleting Messages

You can delete messages by highlighting them and pressing the Delete key. Alternatively, right-click any message (or a selection of them) and select Delete.

The way Evolution handles deleted messages is rather unusual in that messages aren't deleted instantly. They simply disappear from view. You can view all deleted messages by clicking the Trash icon on the left side of the program window, but, technically speaking, the messages haven't been moved to the Trash folder. The Trash folder merely displays in one place messages marked for deletion. You can right-click each message in the Trash view and select Undelete, which will unhide the message, but you can't "empty the trash," as you might be used to doing in other applications.

To get rid of deleted messages permanently, you need to right-click the folder *where you deleted the message* (such as the Inbox) and click Folder ► Expunge. You'll be asked if you really want to delete the messages, as shown in Figure 27-3.

Note If you move any messages from folder to folder, as described later in the "Sorting and Filtering Messages" section, a copy of the mail will end up in the Trash folder. This is because Evolution doesn't literally move messages. Instead, it copies them from the old to the new location and deletes the original. This can be a little disconcerting at first, but there's nothing to worry about. The mail message will remain wherever you moved it, and it won't disappear when you expunge any folders.

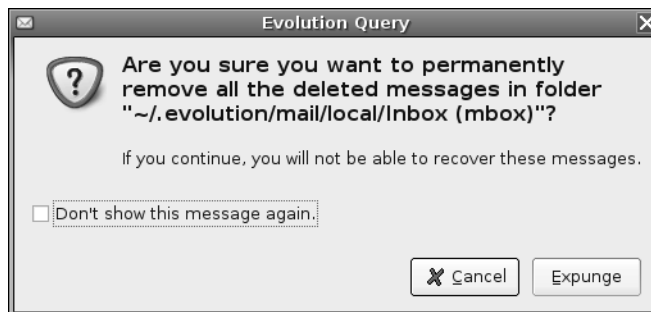


Figure 27-3. To permanently delete messages, it's necessary to expunge them from the folder from which they were deleted.

Flagging Messages

You can flag messages in a variety of ways to help remind you of their status or purpose. The simplest form of flagging is to mark a message as important: right-click the message and select Mark As Important, or click in the space beneath the Important column (this is located to the left of the From column).

Alternatively, you can add several different flags by right-clicking a message and selecting Mark for Follow Up. Your choices range from Do Not Forward to No Response Necessary and Review. This heading will then appear in the message preview at the bottom of the window whenever the mail is selected.

IMPORTING E-MAIL FROM THUNDERBIRD

Back in Chapter 4, I discussed a method of exporting e-mail from various Microsoft e-mail programs, which use proprietary formats, so that it can be imported under Ubuntu. To recap, you can install the Mozilla Thunderbird e-mail client under Windows, import your e-mail into it from Outlook or Outlook Express, and then export Thunderbird's mailbox (.mbox) files for use within Evolution.

If you followed these instructions and now have the .mbox files ready for use with Evolution, it's easy to import them. Click File ► Import. In the Import dialog box, click the Browse button, locate the .mbox file, and click Open. If you have more than one .mbox file, you'll need to import each one manually.

Composing a Message

Creating a new e-mail is as simple as clicking the New button at the top left of Evolution's program window. Fill in the To and Subject details as usual, and then type in the main body of the message.

To add a CC or BCC, click the To: button and select addresses from your contacts list in the dialog box that appears (selecting the CC or BCC button as appropriate). Alternatively, if you would like to have the CC and BCC fields visible and available at all times, click their entries under the View menu of the Compose a Message window.

As with most Microsoft mail programs, new e-mail can be sent either as plain text or as HTML. Plain text mode is the default. To switch to HTML, click the entry on the Format menu. The advantage of HTML mail is that you can vary the style, size, and coloring of text, so can emphasize various words or paragraphs, as illustrated in Figure 27-4. In addition, if you click Insert ► Image, you can insert pictures from the hard disk. Other options on the Insert menu let you insert tables, dividing lines (click the Rule menu entry), and hyperlinks.

The disadvantage of HTML e-mail is that the person receiving the message will need an HTML-compatible e-mail program to be able to read it.

Tip Many people in the Linux community frown on HTML-formatted e-mail and prefer plain text messages.

Words are automatically spell-checked in the new e-mail and are underlined in red if the spell checker thinks they are misspelled. To correct the word, right-click it, click Check Word Spelling, and then select the correctly spelled word from the list.

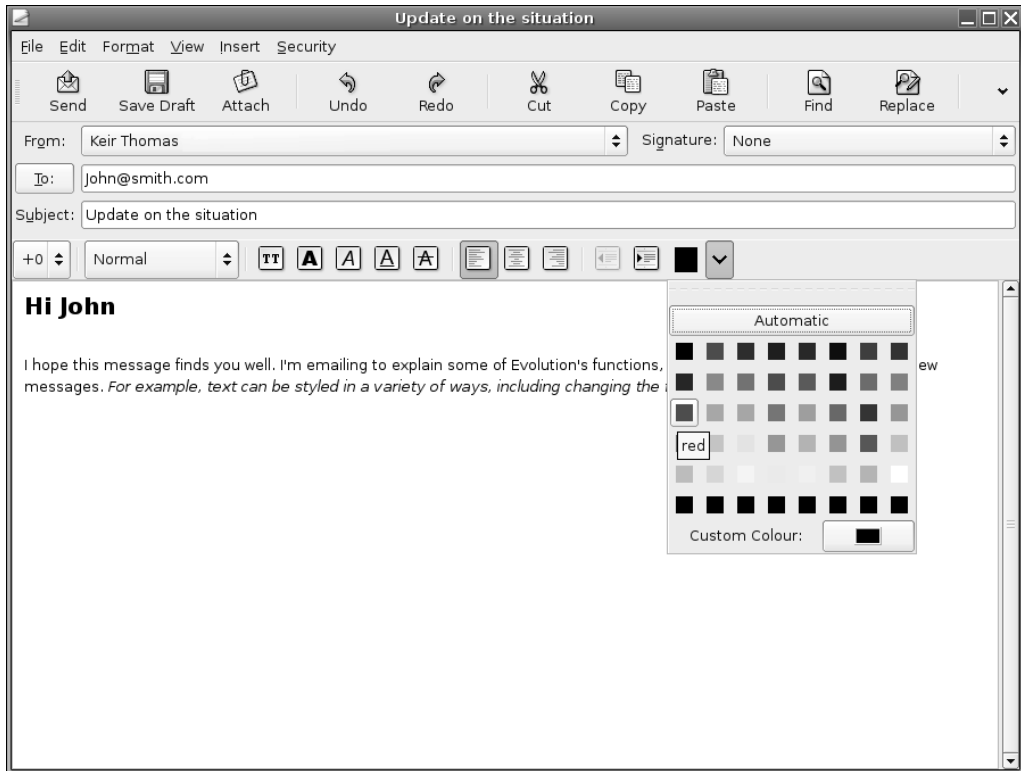


Figure 27-4. New messages can be formatted in HTML, allowing you to change the color of text, as well as its size.

Creating an E-Mail Signature

E-mail signatures are the blocks of text that appear automatically at the end of new e-mail messages you compose. They save you the bother of typing your name and contact details each time. To create an e-mail signature, follow these steps:

1. Click Edit ► Preferences. In the Composer Preferences dialog box, click the Identity tab.
2. Click the Add New Signature button at the bottom right of the window.
3. In the Edit Signature dialog box, type what you wish to appear as your signature. The signature can either be in plain text or HTML (click Format ► HTML). Don't forget that in HTML mode you can insert lines (Insert ► Rule), which can act as a natural divider at the top of your signature to separate it from the body of the e-mail, as shown in the example in Figure 27-5.

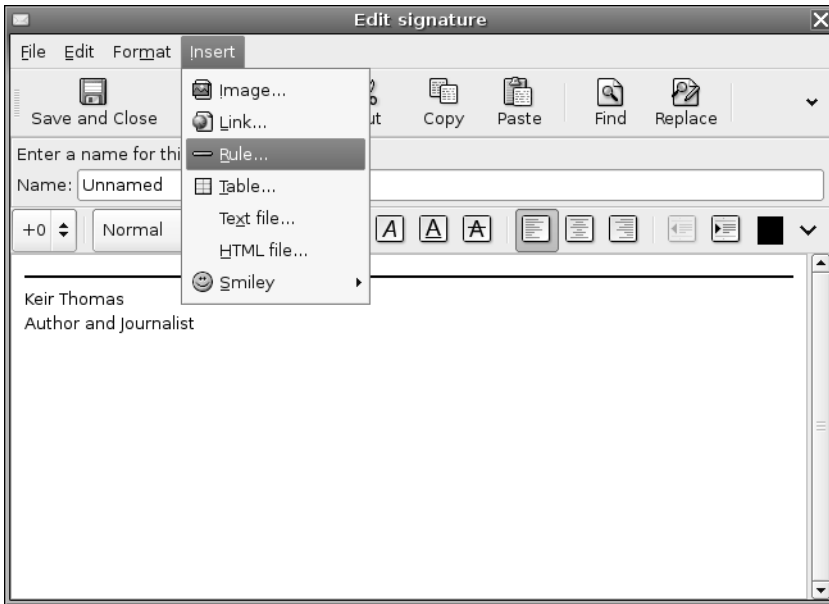


Figure 27-5. Creating an e-mail signature saves you from having to type your contact details each time.

4. Click the Save and Close icon at the top left.
5. Click OK, and then click Close to close.

Tip Enter a few carriage returns at the top of your signature so that, when you create a new e-mail, you have enough space to click and type, without accidentally clicking within the signature.

Advanced E-Mail Tasks

Evolution offers several features that can help keep you to organize your e-mail. You can create new folders, as well as filter, sort, and search through your messages.

Creating New Folders

If you want to better organize your e-mail, you can create your own folders, which will then appear in the list on the left side of the program window.

To create a new top-level folder, which will appear in the list alongside the standard folders (Inbox, Junk, Outbox, and so on), right-click On This Computer and select New Folder. Then make sure that On This Computer is selected in the folder view of the dialog box that appears. Type a name and click Create.

You can also create second-level folders, which will effectively be “inside” other folders and will appear indented below their parent folder within the list. For example, you might want

to create a series of folders within the main Inbox folder to sort your mail from various individuals or organizations. To do this, right-click Inbox, select New Folder, and give the folder a name in the dialog box that appears, as shown in Figure 27-6.

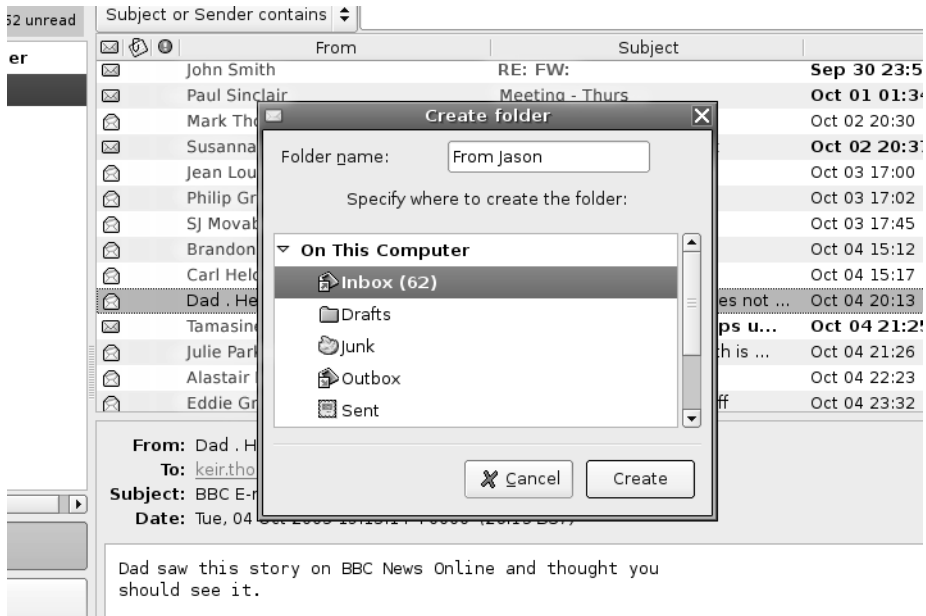


Figure 27-6. You can create your own folders to better organize your mail.

You can then drag-and-drop messages into the new folders, or simply right-click them and select Move to Folder. This can be useful if you wish to select a handful of messages by holding down the Ctrl key. All you need to do then is right-click one of them and select Move to Folder.

You can also copy messages from one location to another, thus producing two copies of the same message. Simply right-click the message and select Copy to Folder. Then select the folder from the list. Alternatively, you can hold down the Ctrl key while you drag the message to the new location.

Dealing with Junk E-Mail

Evolution includes intelligent junk mail filtering. Any mail that Evolution thinks is spam or junk mail will end up in the Junk folder. When you first start using Evolution, you should check the folder regularly, because there's a chance Evolution might have made a mistake. However, this is a good thing, because by right-clicking the message and selecting Mark As Not Junk, the Evolution junk mail filter will be able to better understand what to consider as junk in your particular Inbox.

In a similar way, if you find that Evolution misses a junk e-mail and it ends up in your Inbox, you can right-click it and select Mark As Junk.

To empty the Junk folder, select all the messages (Ctrl+A), right-click, select Delete, and then click Folder ► Expunge. Bear in mind that, as with any folder, once the messages are deleted, they will appear in the Trash view where, if necessary, they can be restored.

Sorting and Filtering Messages

You can filter incoming messages according to practically any criteria, including who sent the message, its subject line, words within the body of the mail, its size, or even if it has attachments. Coupled with the ability to create folders, this allows you to automatically sort messages as soon as they're received.

To set up filters, click Edit ► Message Filters. Click the Add button and, in the Rule Name box, start by giving the new rule a descriptive name by which you'll be able to recognize it in future. You might think this isn't important, but you may create tens, if not hundreds of filters, so being able to identify filters will be very helpful.

As shown in Figure 27-7, the Add Rule dialog box is split into two halves: If and Then. This is rather like a sentence: *If* the selected conditions are met, *Then* the selected actions will take place.

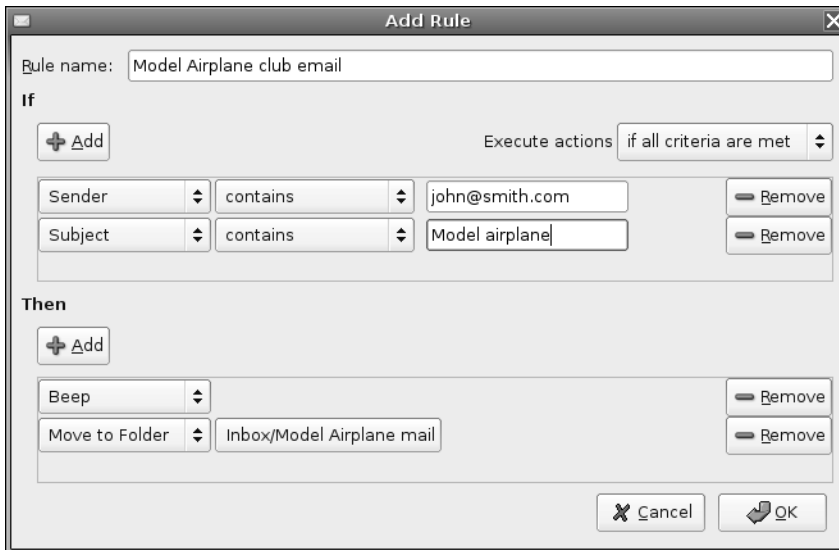


Figure 27-7. Creating message filters lets you automatically organize your e-mail as soon as it's received.

The If part is used to identify the mail. You can select to filter based on almost any criteria, such as who appears in the Sender field of the message, words that appear in the Subject line, the date sent, and so on. Simply select what you require from the drop-down list directly beneath the Add button. In most cases, you'll then need to specify details for the filter. For example, if you select to filter by the address of the individual sending the e-mail, you'll need to provide that e-mail address.

Tip Several If rules can be created. For example, you could create a rule to filter by the address of the sender, and then click the Add button to create another rule to filter by text in the Subject line. By clicking If All Criteria Are Met in the Execute Actions drop-down list, the mail will be filtered only if both conditions are met. By selecting If Any Criteria Are Met from the drop-down list, the mail will be filtered if either condition is met.

Once you've set the If conditions, you need to select from the Then section of the dialog box. This tells Evolution what to do with the filtered mail. The obvious course of action is to move the e-mail to a particular folder, which is the default choice, but you can also delete the e-mail, set a particular flag, beep, or even run a particular program! As with the If rules, you can set more than one condition here, so you can have Evolution beep and then delete the message, for example.

TIPS FOR USING EVOLUTION E-MAIL

In many ways, Evolution is similar to e-mail programs you might have used in the past, but also has a few of its own quirks and idiosyncratic ways of working. Here are a handful of preferences you might want to set to have Evolution behave in a more familiar way:

- **Forward e-mail inline:** If you attempt to forward a message, Evolution will attach it to a new message as a file. The person receiving the e-mail will then need to double-click the file to view the forwarded e-mail, which can be confusing. The solution is to make Evolution forward the message *inline*, which is to say that Evolution will quote it beneath the new mail message, like Microsoft e-mail programs. To do this, click Edit ► Preferences, click Composer Preferences on the left side of the dialog box, click the Forward Style drop-down list, and select Inline.
- **Change the plain text font:** Any messages sent to you in plain text format, rather than HTML, will appear in the message preview pane in a Courier-style font. To have messages display in a more attractive and readable typeface, click Edit ► Preferences, select Mail Preferences on the left side of the dialog box, and then remove the check from Use the Same Fonts As Other Applications. In the Terminal Font drop-down list, select an alternative font. The standard Ubuntu font is called Sans and is a good choice.
- **Always create HTML e-mail:** Evolution defaults to plain text e-mail for any new messages you create. If you want to always create HTML messages, click Edit ► Preferences, click Composer Preferences on the left side of the dialog box, and then put a check alongside Format Messages in HTML.
- **Empty Trash on exit:** To automatically expunge all folders of deleted messages each time you quit Evolution, click Edit ► Preferences, click Mail Preferences on the left side of the program window, and put a check alongside Empty Trash Folders on Exit.

Creating Search Folders

Evolution's search folder feature is a more powerful alternative to message filters. Using search folders, you can filter mail based on a similar set of criteria, but you can choose to include messages in the results that might be *associated* with the filtered messages. For example, if you choose to filter by a specific individual's e-mail address, you can select to have any replies you sent to that person included in the results, rather than simply messages received from her. In addition, you can apply search folders to specific e-mail folders on an ongoing basis, rather than all incoming e-mail.

You can create a new search folder by clicking Edit ► Search Folders. As with creating message filters, clicking the drop-down box beneath the Add button will let you select a criteria by which you can filter. The choices are broadly similar to those for message filters, in that you can filter by e-mail address, size of e-mail, message body, and so on.

In the Include Threads drop-down box, you can select what kind of results you would like the search filter to return:

- None simply returns e-mail messages matching the criteria.
- All Related returns every single message that is associated with the criteria.
- Replies returns results that include replies to the messages returned via the filter.
- Replies with Parents returns results that include replies and also any initial message that you or others might have sent that inspired the message included in the filter results.

Search folder results are listed under the relevant folders on the left side of the Mail mode window.

The search folder feature is very powerful and worth spending some time investigating.

Contacts

Evolution includes a powerful contacts manager component that can catalog information about individuals. At its most basic, the contact manager stores e-mail addresses for use within the e-mail component of Evolution, but you can enter significant additional data about each individual, including addresses, phone numbers, fax numbers, and even a photograph for easy identification. This should allow Evolution to become your sole personal information manager.

To switch to the Contacts mode, click the button at the bottom-left side of the program window. Once in the Contacts mode, you can view information in several ways. Click View ► Current View to choose from the following views:

Address Cards: This is the default view and shows the contacts as virtual index cards arranged alongside each other at the top of the program window. Click the scroll bar beneath the cards to move through them.

Phone List: This shows the contact information as a simple list, arranged vertically, with various elements of the contact's personal information listed alongside, such as phone numbers and e-mail addresses.

By Company: This organizes the data in a similar way to Phone List view but sorted by the company the contacts work for (if such data has been entered into the contact entries).

Adding or Editing Contact Information

By far, the best way of initially building up your contacts list is to right-click e-mail addresses at the head of messages and select Add to Addressbook. This will add a simple contact record consisting of the individual's name and e-mail address.

When using Microsoft mail applications, simply replying to an e-mail from an individual is enough to add that contact to your address book. Evolution is capable of this behavior, too, but the feature isn't activated by default. To set this up, click Edit ► Preferences, click Mail Preferences on the left side of the dialog box, and then put a check in the box marked Automatically Create Entries in the Addressbook When Responding to Mail.

You can then edit the contact details by double-clicking the entry in Contacts mode. This will let you enter a variety of information, as shown in Figure 27-8. To import a photo for this contact, click the top-left icon. You can use any picture here, and you don't need to worry about its size, because it will be resized automatically by Evolution (although its aspect ratio will be preserved). The imported photo will appear on the contact's virtual card.

The screenshot shows the 'Contact Editor' window with the following details:

- Full Name...:** Timothy Jones
- Nickname:** Timothy
- File under:** Jones, Timothy
- Where:** Personal
- Categories:** (empty)
- Email:**
 - Home: tim@example.com
 - Work: tim@example.com
 - Other: (empty)
- Wants to receive HTML mail**
- Telephone:**
 - Business Phone: 01222 234567
 - Mobile Phone: 07999 123456
 - Home Phone: 01222 890123
 - Business Fax: 01222 456789
- Instant Messaging:**
 - AIM: sl1022
 - ICQ: (empty)
 - Yahoo: (empty)
 - GroupWise: (empty)

Figure 27-8. A lot of information can be entered for each contact and, by clicking the button at the top left, you can also add a photograph.

Creating a Contact List

Contact lists are simply lists of e-mail addresses. Once a list is created, you can right-click its entry in the contacts list, and then choose to send a message to the list or forward it to someone else as a vCard. The obvious use of contact lists is for sending group e-mail messages.

Note A vCard is a virtual business card. Effectively, it's a small file that contains personal information. As well as personal data, vCards can contain pictures and audio clips. They're understood by practically all business-level e-mail programs, including Microsoft Outlook and Apple Mail.

To create a contact list, click the small down arrow next to the New button in Contacts mode, and select the option from the list. Then simply click and drag contacts from the main program pane onto the bottom of the Contact List Editor pane. This will automatically add their names and e-mail addresses. Alternatively, you can type their e-mail addresses manually into the Members field, and then click the Add button, which can be useful if the individual isn't in your contact list.

By checking the Hide Addresses When Sending Mail to This List option, you can ensure that the e-mail addresses are added to the BCC field of a new message, so people on the list don't see the others on the list.

Calendars

The Calendars mode of Evolution allows you to keep an appointments diary. Entries can be added in half-hour increments to the working day, and you can easily add events to days that are weeks, months, or even years in advance. Viewing a day's appointments is as simple as clicking its entry in the monthly view at the top right of the program window.

Specifying Appointment Types

You can make the following three types of diary entries:

Appointments: These are events in your diary that apply to you only. You might have a meeting with an external supplier, for example, or might simply want to add a note to your diary to remind you of a particular fact.

All Day Events: These are appointments that take the entire day. For example, a training day could be entered as an all day event. However, all day events don't block your diary, and you can still add individual appointments (after all, just because your day is taken up with an event, it doesn't mean you won't need to make individual appointments during the event). All day events appear as a blue bar at the top of the day's entry in your diary.

Meetings: Meetings are like appointments but you also have the option of inviting others to attend. The invitations are sent as iCal attachments to e-mail, so users of Microsoft Outlook should be able to reply to them (provided Outlook is properly configured; see the program's documentation for details, and note that iCal is sometimes referred to as RFC 2446/2447). Once an individual receives a meeting invitation, he can click to accept or decline. Once Evolution receives this response, the individual's acceptance or declination will be automatically added to the diary entry.

Adding or Editing a Diary Entry

To add a new diary entry, simply select the day in the monthly view at the top right, and then select the time the appointment is to start. Then right-click and choose either an appointment, all day event, or meeting. To edit an already existing diary entry, double-click its entry in the list.

At its most basic, all an appointment needs in order to be entered into your diary is some text in the Summary field, as shown in Figure 27-9. By default, appointments and meetings are assumed to last for half an hour, but you can change this by clicking down arrow in the End Time field. For what it's worth, appointments can go on for days—just select a different date in the End Time field!

The screenshot shows a dialog box titled "Meeting - Meeting with suppliers". It has four tabs: "Appointment", "Recurrence", "Scheduling", and "Invitations". The "Appointment" tab is selected. Under the "Basics" section, there are fields for "Summary" (containing "Meeting with suppliers"), "Location" (containing "NY Office"), "Classification" (a dropdown menu set to "Public"), and "Calendar" (a dropdown menu set to "Personal"). There is also a "Categories..." button and a "Description" text area containing "Meeting with suppliers to discuss contracting issues and resolve performance concerns." The "Date and Time" section has "Start time" (29/04/2006, 13:30) and "End time" (29/04/2006, 16:30) pickers. There are checkboxes for "All day event" (unchecked) and "Show time as busy" (checked). An "Alarm" section is set to "15 minutes before appointment" with a "Customise..." button. At the bottom, there is an "Attachment Bar (drop attachments here)" and buttons for "Help", "Cancel", and "OK".

Figure 27-9. When creating a new appointment you can add all the details you need, but don't forget to set the end time!

By clicking the Recurrence tab, you can set the appointment to be booked into your diary according to certain intervals. Start by putting a check in the This Appointment Recurs box, and then select a time interval. For example, selecting 1 week will mean that the appointment is booked into your diary automatically on a weekly basis. After this, select a day of the week for the recurring appointment. Following this, you must either specify the number of recurrences or

simply select Forever from the drop-down list. Then click the Add button to add the details of the recurring event to the appointment.

In the case of meeting appointments, you can click the Invitations tab to invite others to the meeting, via iCal invitations which will be sent out by e-mail as soon as you've finished creating the appointment. Simply click the Add button and then, in the empty field that appears, start typing the contact name of the individual you want to invite. If the person is already in your contacts list, the name will be automatically completed, but you can also type individual e-mail addresses.

Clicking the Scheduling tab will show you who can and can't attend, according to replies to the invitations sent out (obviously, this is a feature you'll be using after you initially created the appointment). On the left side of the dialog box, you will see the list of attendees and also their status: whether they've accepted or not, or whether they've sent a busy/tentative reply (in which case, you might choose to reschedule the meeting).

Tasks

The Tasks mode is the simplest component within Evolution and allows you to create a to-do list. After you've made an entry, clicking the check box alongside it will mark it as completed. Completed items appear with strike-through, as shown in Figure 27-10.

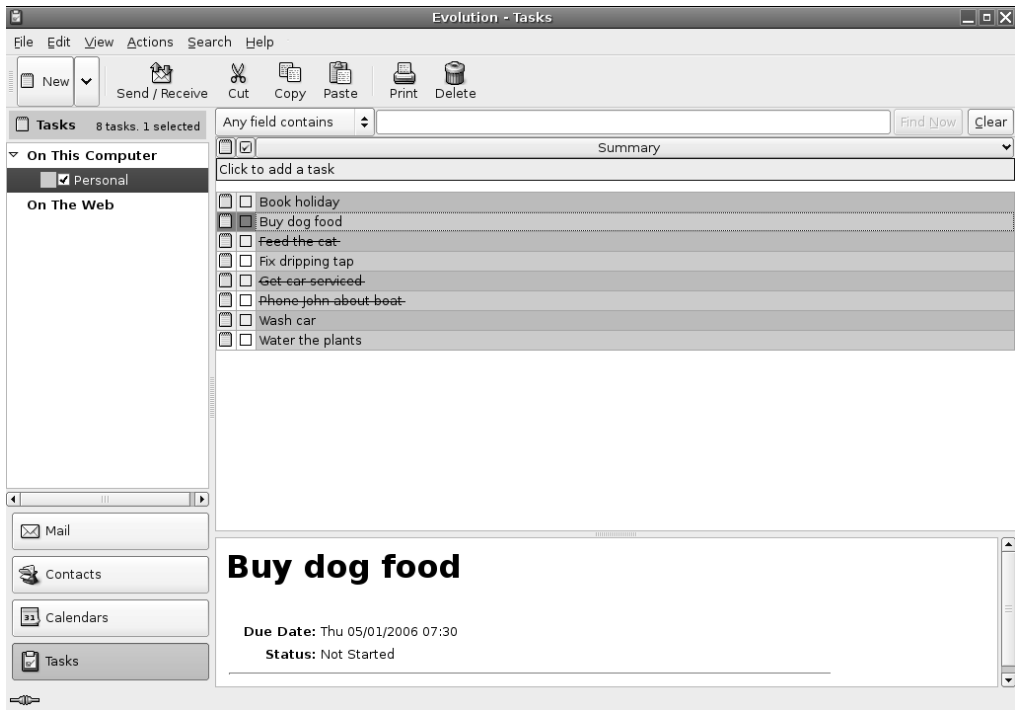


Figure 27-10. *The Tasks mode lets you catalog chores that you have to do during the day.*

Switch to Tasks mode by clicking its button on the bottom-left side of the program window. To add a new task, click the bar that reads Click to Add a New Task. Type a description of the task, and then press Enter. You will then be able to enter further tasks in the same field.

Double-clicking a task allows you to fine-tune its details. For example, you can add a due date, so you'll know when the task must be completed. You can also add a description for future reference.

By clicking the Status tab, you can also set a percentage figure for completion of the task, as well as its priority, ranging from Low to High.

Summary

This chapter has been a whistle-stop tour of Evolution's main features. We've looked at e-mail creation and organization, contacts management, working with the appointments calendar, and editing the task list.

Evolution is a powerful program. Be sure to take a look at its help documentation (Help ► Contents) to learn more about it.

In the next chapter, we'll look at running Microsoft Office under Ubuntu. Yes, you read that right. It's entirely possible to run certain Windows applications under Linux.



Running Microsoft Office Under Ubuntu

The title of this chapter might sound somewhat strange. How can a Windows program be run under Linux? The answer comes courtesy of the Wine project and the set of add-on programs it offers for Linux. Wine is an acronym for Wine Is Not an Emulator, which is a way of saying that the Wine software doesn't emulate Windows but instead re-creates certain parts of it called the Application Programming Interface (API).

In theory, the vast majority of Windows program can be run using Wine. The reality is that some work easily, some require a little configuration to run, and others don't work at all.

Getting a Windows program to run under Wine is rarely difficult, but essentially it's a "hack"—a way of making something work when it's not intended to do so. Such hacks are the lifeblood of Linux, but this should serve as a warning to the less-experienced Ubuntu user: don't expect to be able to simply insert the Microsoft Office CD, click a few times, and have a working Office installation! A handful of preparatory steps are necessary.

In this chapter, you'll learn how to install both Wine and Microsoft Office. Because the presence of Internet Explorer is critical to many pieces of Windows software nowadays, we also cover the installation of this web browser.

Installing and Configuring Wine

Wine isn't installed by default under Ubuntu. Although you usually should download new software from the Ubuntu repositories, in the case of Wine, it's a better idea to download it directly from the people behind the software. This is because Wine is constantly being updated and made more compatible with new Windows software, and you'll find the latest version at the SourceForge.net site.

Downloading Wine

To get the latest version of Wine, add a new download repository to the Synaptic Package Manager, as follows (see Chapter 8 for details on setting up the Synaptic Package Manager):

1. Select System ► Administration ► Synaptic Package Manager. Click Settings ► Repositories.
2. Click the Add button, and then the Custom button in the dialog box that appears. In the APT Line field, enter the following (note the space inserted before binary/), as shown in Figure 28-1:

```
deb http://wine.sourceforge.net/apt/ binary/
```



Figure 28-1. Getting the latest version of Wine ensures maximum compatibility with Windows programs.

3. Click the Add Repository button.
4. Click OK in the Software Preferences dialog box, and then agree to the reloading of package listings.
5. In the main Synaptic Package Manager window, click Search and enter wine as a search term.
6. In the list of results, look for the wine entry (this will simply be the word wine on its own, without any additional text). Click the check box, and then click Mark for Installation.

7. Search for `winetools` and mark it for installation.
8. Click `Apply`. You'll be warned that some of the downloads you're about to install are not authenticated. However, in this case, you're downloading from a trusted location (<http://sourceforge.net>, in fact, the home of many key open-source projects). So it's safe to click `Apply`.
9. When Wine has downloaded and installed (it was around 15MB on my system), close the Synaptic Package Manager.

CROSSOVER OFFICE

The developers behind Wine also produce a proprietary version of Wine called CrossOver Office. Technically speaking, this is actually no different from any other version of Wine, except that a GUI-based front end is added to make installation of programs simpler. It also includes several scripts to overcome various incompatibilities between Windows programs and Wine. This makes it possible to install the very latest version of Office, for example. In addition, CrossOver Office also lets you use Internet Explorer-based plug-ins within Linux-based browsers.

You can learn more about CrossOver Office by visiting its official web site at www.codeweavers.com. CrossOver Office is available for a fee from CodeWeavers, but considering they are one of the major sponsors of the Wine project, this is an excellent way of investing in the community.

Configuring Wine

The Sidenet Wine Configuration Utility is a series of scripts by a third-party programmer (who isn't related to the Wine project) that automatically configures Wine so that it's much more compatible with current software. Sidenet automates the installation of Internet Explorer 6 and Windows Media Player 6.4 (you might find the versions numbers differ). It also tweaks the Wine configuration in a handful of ways to make it work with a lot of software.

Note A script is a series of commands within a text file. By running the script, you work through the series of commands automatically.

The Sidenet script can create several different types of Wine configuration. The first is a completely Microsoft-free Wine installation, which means that it won't install any Windows-based system files or Internet Explorer. The other options let you install Internet Explorer and Windows Media Player 6. This is strongly advisable, because many Windows programs demand the presence of these two programs (particularly Internet Explorer). Additionally, it's very useful to download several additional components that can then be installed using Sidenet.

Caution According to the wording of Microsoft's End User License Agreements, you can download and install the Windows components only if you have a valid Windows license. That's to say you can download Internet Explorer and Windows Media Player only if you own Windows XP, 2000, Me, 98, or 95. If you don't own these operating systems, then you'll be violating copyright laws. Also, you can install DCOM98, an additional component, only if you have a Windows 98 license; to do otherwise is to break copyright laws.

Here are the steps for installation:

1. Open a web browser and go to <http://sidenet.ddo.jp/winetips/config.html>. From there, download the latest version of the Sidenet binary release.
2. When the download of Sidenet has finished, open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type the following at the prompt (change the filename listed on the tar line to match the name of the file you downloaded; these steps assume the file was downloaded to the desktop):

```
cd ~/Desktop
tar xzf wine-config-sidenet-1.9.0.tgz
```

This will switch to where the download has been saved and unpack the downloaded archive into its own directory.

3. Download the following manually and place the downloaded files in the Sidenet directory you created earlier (on my system, this was wine-config-sidenet):

- DCOM98: www.microsoft.com/com/default.mspx
- Windows Installer add-in: www.microsoft.com/downloads/release.asp?releaseid=32831
- MFC 4.0 Runtime: www.softlookup.com/download.asp?ID=10315

The MFC 4.0 Runtime is supplied in a zip file. Open it with File Roller and extract all the contents of the zip to the Sidenet directory.

Note The URLs for the additional Windows components were correct as this book went to press but it's possible that they may change. If that happens, searching the Web for the files in question should return ample download sites.

4. Once all the files are downloaded, click the GNOME Terminal window and type the following (this again assumes that you created the Sidenet directory in your desktop directory):

```
cd ~/Desktop/wine-config-sidenet
./setup
```

5. Answer the prompts until you are offered the choice of the kind of Wine setup you want. Choose to manually install (this was option 3 on the version of Sidenet that I installed, but might be different for you).
6. You need to answer more questions about which programs you want to install. Answer Yes to each, and agree to create the C:\ folder in your /home directory.

Note During installation of Sidenet, the Windows components you install might ask you to reboot. It isn't necessary for you to actually reboot Linux. Instead, a Windows reboot is faked within the Wine software (this is done automatically by Sidenet). You can do this yourself at any time from the GNOME Terminal window by typing `wineboot`.

7. Eventually, the DCOM98 installation window will appear. As you would within Windows, work through the options, agreeing to install the software.
8. Following this, the Internet Explorer installation program will download and then start. Again, work through the installer. Choose the Install Now ► Typical Set of Components option. The program will load, as shown in Figure 28-2.

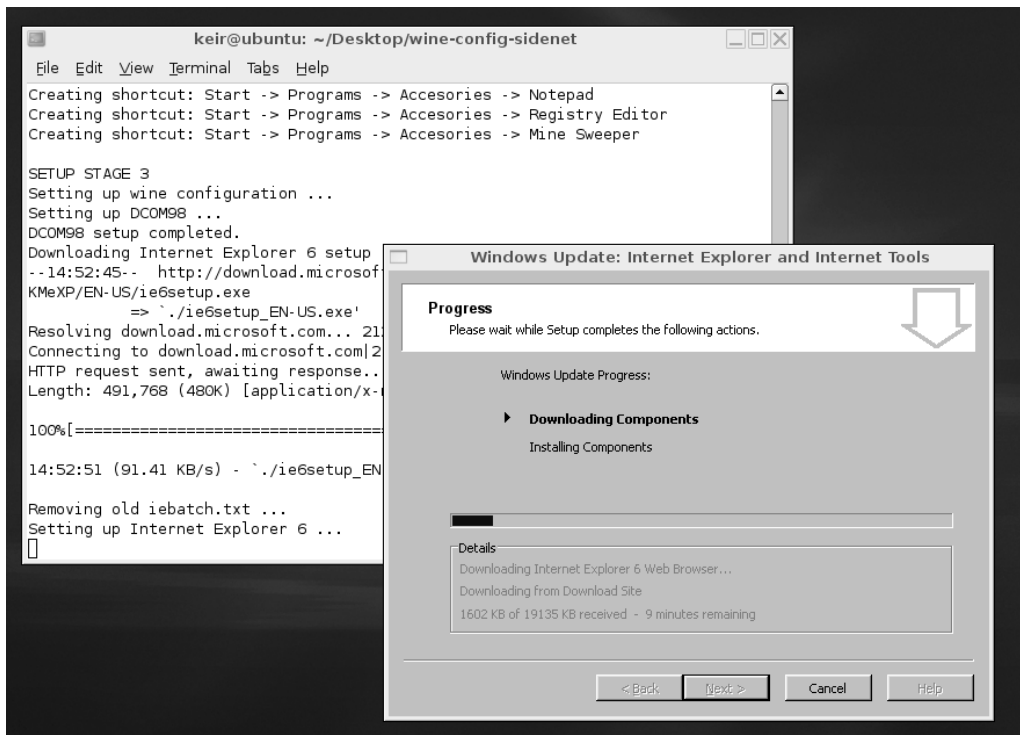


Figure 28-2. The Sidenet script will download and install Internet Explorer and also configure vital Wine settings.

9. At the end of installation, an Internet Explorer 6 window will open, informing you that the installation has finished. Closing this window will also end the script.
10. You can delete the `wine-config-sidenet` folder, as well as any other files you downloaded.

Note I saw an error dialog box about `ieinfo5.ocx` when installing Internet Explorer. Don't worry if you see the same thing, or even a handful of other errors. Chances are everything will still run smoothly.

Showing Hidden Files on CDs

You need to make one final change so that Wine can see hidden files on certain Windows CDs. This involves manually editing `/etc/fstab`, the configuration file that controls how and where file systems are mounted (`fstab` and mounting are covered in Chapter 15), to add the word `unhide` to the line relating to the CD-ROM within `fstab`. Here's the procedure:

1. Make sure the CD-ROM drive is unmounted. If you're browsing the CD, close the file window and make sure that any programs that might access the CD are closed, too.
2. Type the following at the GNOME Terminal command prompt:

```
umount /media/cdrom
```

3. Load the `fstab` file into the Gedit text editor:

```
sudo gedit /etc/fstab
```

4. Search for the line that contains the words `/media/cdrom0` and add `unhide` just before the `user,noauto` entries in the `Options` column. The final line should look something like this:

```
/dev/hdc /media/cdrom0 udf,iso9660 unhide,user,noauto 0 0
```

5. Save the file and quit Gedit.

Installing Microsoft Office

The best choice of Microsoft Office product to install is Office 2000. This has a proven record of installing and generally running well under Wine. It also uses file formats that are, for all practical purposes, unchanged in the more recent versions of Office (XP and 2003) and universally recognized in the business world.

Note The Wine programmers are constantly striving to make Wine compatible with all Windows software packages, but at the time of writing, Office XP and 2003 didn't work, although some people have reported moderate success. The Wine project has a database explaining which software packages work and which don't. Often, any tricks or tips others have used to make the software work are also detailed. You'll find the database at <http://appdb.winehq.org>.

Follow these steps to install Office 2000:

1. Insert the Office 2000 CD, and then click on the GNOME Terminal window.
2. Mount the CD, switch to the `/media/cdrom` folder, and use Wine to run the Office setup program:

```
mount /media/cdrom
cd /media/cdrom
wine SETUP.EXE
```

3. Work your way through the Office 2000 installation program as you would if you were running Windows. Select the default installation option.

Installation didn't run perfectly smoothly on my test PC, although it may work fine for you. I received several errors messages, which I simply clicked through. At the end, the Office installation program seemed to be locked into a cycle of configuring itself. Eventually, it finished and presented the Office installer repair dialog box. I clicked to cancel this, and was then told that installation had failed. However, I found that Office 2000 had, in fact, installed perfectly.

Another test I undertook resulted in error messages about Internet Explorer being split from the Windows Update process. In addition, the installation window was very small and the text within it largely illegible. However, I was able to work through the installation wizard by pressing Enter a few times. You may find it worthwhile to install the Microsoft Core Fonts package before installing Office; see Chapter 21.

As I mentioned at the beginning of this chapter, Wine is something of a hack, so you may encounter some challenges. However, with a little common sense, you should be able to get Office 2000 onto your computer.

Running Office Components

In order to install programs, Wine creates an entire pseudo file system that mirrors a typical Windows hard disk. This is normally located in the hidden `.wine` directory in your system folder, but the Sidenet script you used to configure Wine will have placed it in your `/home` directory. Therefore, you should be able to browse to `/home/<username>/c/Program Files` to see a listing of all installed Windows software (replacing `<username>` with your username).

Running the Office programs is a matter of navigating to the `/home/<username>/c/Program Files/Microsoft Office/Office` folder and running the executables via Wine. For example, to run Microsoft Word on my system, I typed:

```
wine /home/keir/c/Program Files/Microsoft Office/WINWORD.EXE
```

To run Excel, you can type `EXCEL.EXE`. PowerPoint is `POWERPNT.EXE`, Access is `MSACCESS.EXE`, and Outlook is `OUTLOOK.EXE`. Don't forget that these must be typed in uppercase, because uppercase and lowercase matters in Linux. Figure 28-3 shows an example of running Word under Ubuntu.

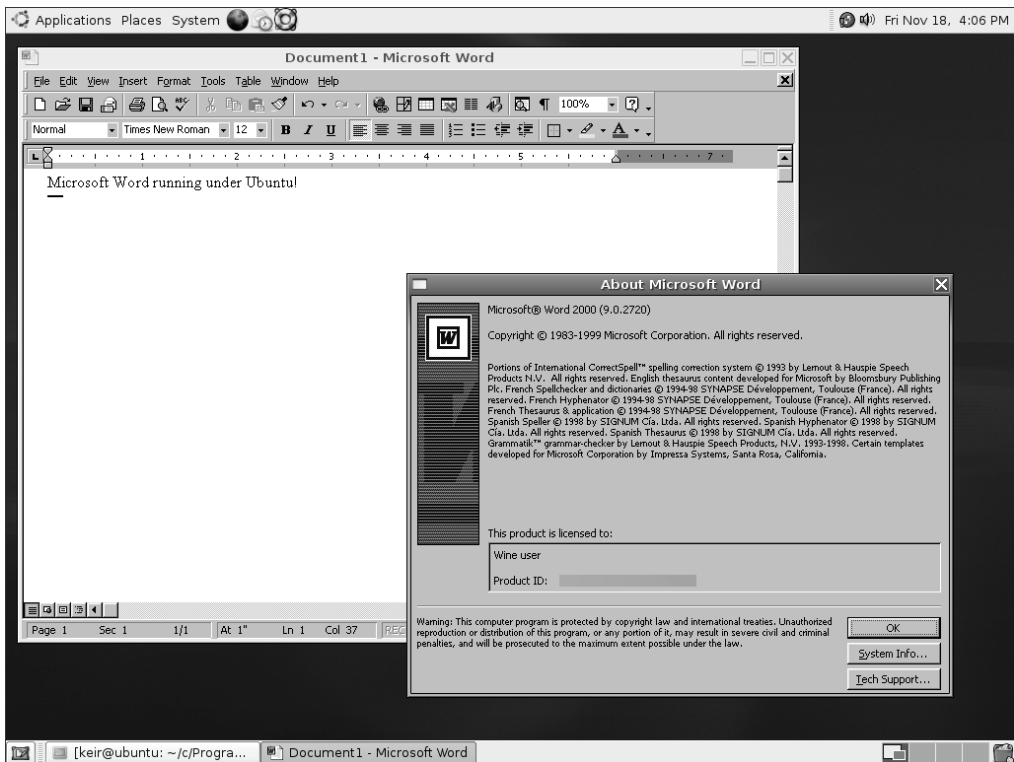


Figure 28-3. It's entirely possible to run Microsoft Word under Ubuntu. Just add a little Wine!

Alternatively, you can easily create desktop shortcuts for your Office programs. See Chapter 10 for details on creating shortcuts.

Aside from being a subdirectory within it, the dummy `C:\` drive is entirely separate from your main file system, which is to say that no Windows programs will ever be installed outside this subdirectory. This should give you some peace of mind when it comes to threats from viruses and worms. In theory at least, only the files in this directory will be affected should you pick up an infection. (The Wine developers have theorized that a Windows virus running under Wine could infect your main system, but it would have to be created with knowledge that it might be running under Wine, and no such viruses have yet been found.)

Note I experienced a few weird bugs when running some Office programs. For example, the text in the Office Assistant dialog boxes didn't appear. I also found that the clipboard didn't work correctly when two or more Office applications were running (it seemed fine with just one application). You might encounter similar problems, but you should find that the key functions work.

Installing Other Windows Programs

A lot of Windows applications will install and work under Wine. Installing them is usually a matter of typing this:

```
wine <installer_name.exe>
```

As I mentioned earlier, the Wine database (<http://appdb.winehq.org>) often details the best way to get Windows software working. Very nearly every mainstream Windows program can be made to work, although it may take a little extra effort. If there's a program that doesn't work, you can suggest it to the developers. The site contains a list of the top 25 programs that don't work, but which people would like to see up and running under Wine.

In my tests, I managed to use Wine to install the popular WinZip archiving application, as well as the Winamp Lite audio software and Adobe Photoshop 6. Figure 28-4 shows each up and running.




Figure 28-3. Although it sometimes takes a bit of effort, virtually any Windows application can be made to run under Ubuntu.

Summary

In this chapter, we looked at how you can use the Wine project to install and run various Windows programs under Ubuntu, including Microsoft Office. Because using Wine is not without problems, we looked at how it should be configured, and then walked through the steps needed to install Microsoft Office 2000. Then you saw how to install other Windows programs.

In Part 7 of the book, starting with the next chapter, we look at the techniques you need to know to keep your Ubuntu system running smoothly. Chapter 29 explains how to install Linux software.

PART 7



Keeping Your System Running



Installing Software

One of the fun aspects of running any operating system is the ability to expand it—to add in new software over time to improve your workflow or just for entertainment value.

Linux is blessed in this regard, because tens of thousands of software titles are available to meet just about every need. However, if you've tracked down the ideal software title, there's just one barrier to overcome: actually installing it on your system.

Installing software under Ubuntu isn't the same as with Windows. Users are afforded a lot more power over what happens to their system, but this comes at the expense of needing to take a little time to understand the terminology and techniques.

In this chapter, you'll learn the basics of software installation under Ubuntu. Then we'll take a closer look at how the Synaptic Package Manager works. Next, you'll learn how to install software from the command prompt. Finally, we'll cover installing from source.

Software Installation Basics

Installing programs on Windows is relatively easy. If you wish to use the Winamp media player, for example, you can browse to the web site, download the installer .exe file, and install the software. Although you might not realize it, a lot of work goes into making this seemingly simple task possible. Once the original software has been created by the programmers, it must be made into a form that you, the end user, can use.

The first thing to happen is that the software is *compiled*. This is the process of turning the source code created by programmers into an actual file (or set of files) that can be used on a daily basis. On most systems, compiling source code involves a lot of number crunching. This takes time—whole days, in some cases.

Once the program files have been compiled, there needs to be a way they can be installed on various systems and easily transported across the Internet. This is where *packaging* comes into the equation. Programs usually consist of many files. To make each program file individually available would present the risk of some getting lost or corrupted, and the program wouldn't work. Therefore, the files are usually combined into a single archive file, or *package*. In addition, third-party system files are added to ensure compatibility on all computers and an extra program, called an *installer*, is added so that users can quickly get the files onto their system.

All of this means that, to be able to install a program like Winamp on Windows, you just need to download the installer .exe file and run it once. No more work is necessary.

Linux is a little more involved, largely because it never assumes that users want their environment to be simplistic and with limited options. However, most Linux distributions still embrace the paradigm of packaging installation files.

Formats of Linux Installation Files

If you visit the web site of a particular Linux application, you may find that the application is available to download in a number of different formats. The program will almost certainly be available as source code—the original listing that the developer created. It might also be available as a binary, or as a package file.

Note Linux isn't the only operating system for which open-source programs are created and used. There are open-source projects for both Windows and Apple Macintosh, many of which are hosted at the <http://sourceforge.net> web site. Many other less widely used operating systems also rely on open-source software to a greater or lesser extent.

Here are the formats by which Linux software is normally distributed:

Source code: Programmers write their software in various programming languages, such as C and C++, and the code that results is known as *source code*. To make source code usable, it must be *compiled* into a *binary file*. Because the cornerstone of the Linux philosophy is the sharing of source code, you'll almost always find the source code of a program available at the developer's web site. You can then download and compile this on your own system (or, if you're so inclined, study the source code to further your understanding). Although compiling source code isn't very hard to do, it's more convenient to download either a binary version of the program or a package.

Binary files: You might find ready-made binary files are available at the developer's web site. In other words, the programmer has, as a service to users of the program, compiled the source code so that it's ready for use as soon as it's downloaded. Sometimes, binary files come with scripts to help you install them. However, in most cases, you simply place the files in a convenient location on your hard disk, and then run them from there.

Note In both the case of source code and binary files, the files usually come in a *tarball*, which is a single archive file containing other files. A tarball isn't by definition compressed, but usually the `bzip2` or `gzip` tool is used to shrink the file to ease transportation across the Internet.

Self-installing binaries: Some larger programs are made available as self-installing binary files. This comes very close to the way Windows works because, when the file is executed, a GUI-based wizard takes you through installation. If you download OpenOffice.org from the official web site (www.openoffice.org), for example, you'll end up with a single 80MB+ file, which you then simply run from the command line to install the program.

Package files: In many cases, you'll find that a package file of the program is available. In this case, someone has compiled the software files and put them altogether in a single, easily transportable file. Ubuntu package files end with `.deb` file extensions, but other Linux distributions use other package formats, such as RPM (Red Hat Package Manager). The software that oversees packages is called a *package manager*.

Note As a blanket rule, an installation package created for one distribution won't be compatible with another. It's possible to use a program called `alien` under Ubuntu, which aims to convert packages between distributions and different package formats, but this should be done as a last resort. The first resort is to obtain a package specifically designed for your Linux distribution.

Packages Explained

Of all the formats for Linux installation files, packages are by far the most common and popular. This is because a well-implemented package-management system is able to install programs, upgrade them, and uninstall them, all with just a few keystrokes or clicks of the mouse. This approach vastly reduces the amount of work required to get new software onto your system.

The package-management system used under Ubuntu is borrowed from the Debian Project, where it was never given a proper name. Most people refer to it simply as the *Debian package-management system*. Ubuntu relies on packages in order to install software, and the Ubuntu CD contains hundreds of them, as shown in Figure 9-1.

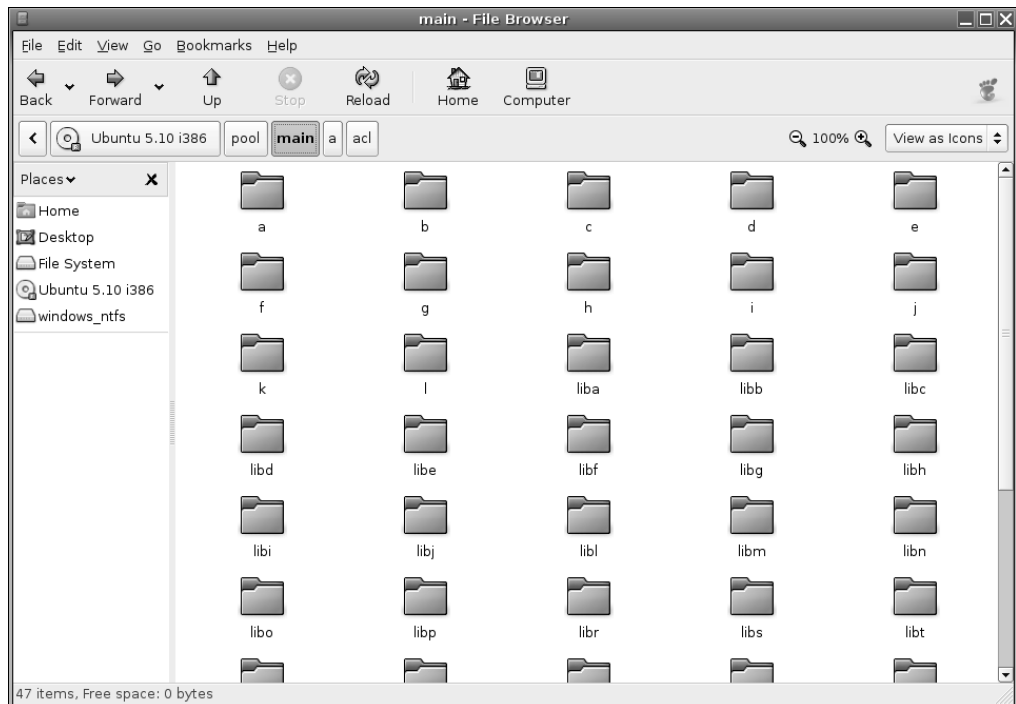


Figure 9-1. There are many packages under Ubuntu.

It's important to understand what a package actually is and what it contains. With Windows, an installation `.exe` file is effectively a piece of software combined with an archive of files. When you run the executable, it triggers a small program within the executable that then unpacks the contents of the file and installs them to the hard disk.

In contrast, Debian packages (.deb files) merely contain the program files, along with a handful of configuration files to ensure the software is both installed and set up correctly. Debian package files are useless without the various pieces of software that are used to manipulate them and that constitute the Debian package-management system.

This approach has a number of benefits. The package-management software builds its own database, so it knows exactly which programs are installed at any one time. It also keeps tracks of version numbers. This gives users much more control over the software on their system.

It also means that if a program starts to act strangely, its configuration files can simply be refreshed using the package manager. There's no need to uninstall and then reinstall the software, as is so often the case with Windows programs.

SOFTWARE VERSIONS

Because most Linux software is open source, a curious thing happens when it comes to software versions. Rather than there being just one "official" version of a program, such as with most Windows software (where you must download the official version of the file), many individuals and organizations take the source code, compile it, and make their own package files available for others to use.

For example, virtually all the software installed with Ubuntu has been compiled by Ubuntu developers. This means it can be quite different from what's "officially" available at the programmer's web site. In some cases, the source code is tweaked, so that notorious bugs are fixed, or a different look and feel applied to the software, so it integrates with the distribution. Often, the configuration files are changed so that the software works properly under Ubuntu, such as integrating with other software packages.

The programmer behind the software doesn't mind when such things happen, because this way of working is part and parcel of open-source software. In fact, the programmer is likely to encourage such tweaking.

Because of this, the first place to look if you want any additional software is not the developer's web site, but the Ubuntu online software repositories. This way, you'll get an officially sanctioned Ubuntu release that will fit in with the rest of your system and won't require much, if any, additional work to get it up and running.

Dependency Management

In addition to actually installing programs, a vital component of package-management software is taking care of *dependencies*. Put simply, the package manager must ensure that if you install a piece of software, any software it relies on to work properly is either already present on the system or is installed at the same time.

Sometimes, the software you want to install might rely on other programs on your system, but more often, the dependencies take the form of system libraries.

It helps if you realize that not all packages contain software that you, as a user, will use directly. Some packages contain nothing but library files, which are shared pieces of code that are equivalent to .dll files under Windows. The key library on an Ubuntu system is the GNU C Library, without which the Linux kernel couldn't function, and this is provided by the lib6c package. But practically every program has its own needs when it comes to library files, and these requirements must be handled by the package manager.

Note One reason why Windows installation files are often so large is that they typically come with all the system files they need in case those files are not already present on the system. Windows isn't alone in this regard; installation files for the Apple Macintosh are similar.

Dependency management, sometimes referred to as *dependency resolution*, is a key feature of the Debian package-management system. You might already have seen this in action: when you mark a program for installation in the Synaptic Package Manager, you may be informed that additional packages need to be installed. If you refuse this request, the program will be deselected. The package manager simply won't let you install a piece of software unless you let it take care of the dependencies. There's no point in letting you install the software because it probably won't work.

But dependency management doesn't just mean adding in packages that a piece of software needs. It might also mean removing packages already present on your system. This might happen if other packages are incompatible with the software you want to install, something that's referred to as *package conflict resolution*. In addition, sometimes you might want to remove a package that other packages rely on, a situation known as *reverse dependency*. In such a case, the package manager either prevents you from removing that software or removes the reverse-dependency packages, too.

Note You can see what dependencies a package has by right-clicking its entry in the Synaptic Package Manager, clicking Properties, and then clicking the Dependencies tab.

APT and Package Repositories

Installing, removing, and otherwise manipulating packages under Ubuntu is handled by the `dpkg` software. However, `dpkg` is not too capable when it comes to managing dependency resolution. It's a fairly basic program whose primary function is to handle the minutiae of software installation, such as getting programs on and off the system.

Because of this, the Advanced Packaging Tool (APT) was created. If you use APT to install a piece of software, you'll be told about any dependencies it needs. Not only that, but the required packages will be automatically added to the list of software that's to be installed.

Note The Synaptic Package Manager is simply a GUI front-end for the APT system. You can see this clearly when you're installing or removing a package. In the Apply Changes dialog box that appears after software has downloaded, click Terminal, and you'll see the output of various APT commands.

APT is able to automatically install the extra software because it's built on the principle of software repositories. These repositories are collections of software arranged into libraries from which you, as a user, can retrieve what you want. Whereas `dpkg` is used to install files

you've downloaded or otherwise got onto your hard disk, the APT tools automatically grab and install packages from repositories. Software repositories can exist on a hard disk or a CD/DVD (the Ubuntu installation CD is a relatively small APT repository), but APT is most frequently used to access online software repositories.

Back in Chapter 8, you set up online software repositories when you configured the Synaptic Package Manager, as shown in Figure 29-2. This enabled you to install software later, as you've done throughout this book. What you actually did was to configure APT's `sources.list` file, found at `/etc/apt/sources.list`. This file contains details of where APT can find the repositories, whether they are online or anywhere else, such as the installation CD.

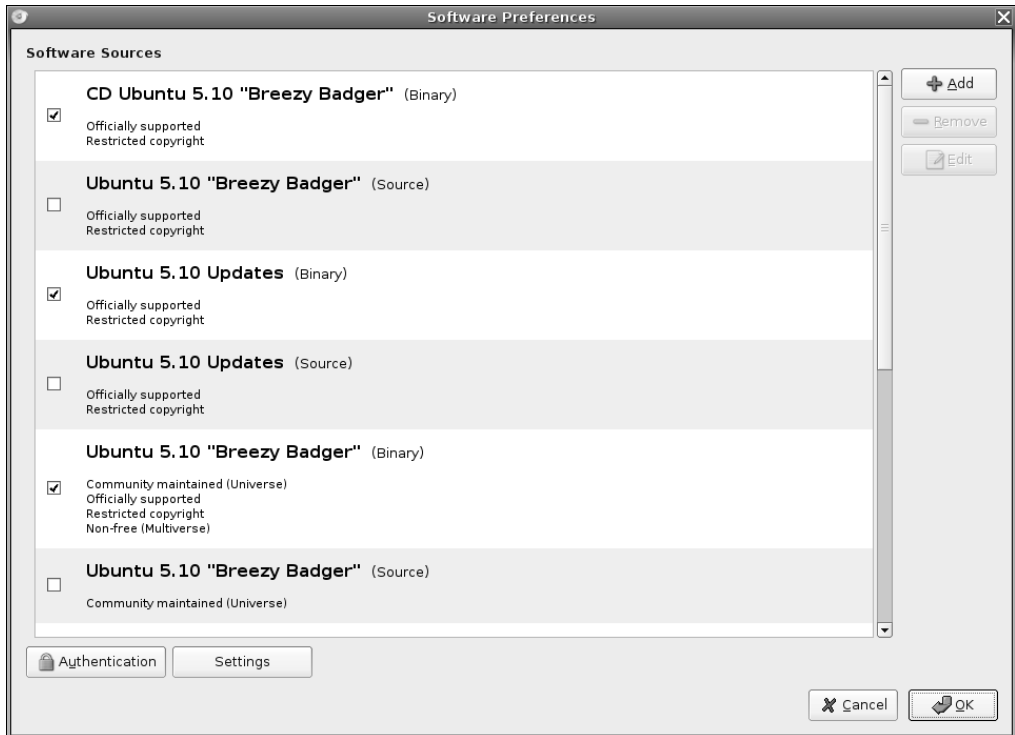


Figure 29-2. Synaptic needs to know where the software repositories are located, regardless of whether they are on CD or online.

Categories of Repositories

Regardless of whether they're online or on a CD, Ubuntu repositories are strictly categorized according to the type of software they contain:

Main Distribution: This repository contains the packages that are used to install Ubuntu. It's a little like the installation CD's repository, but online and containing additional packages that wouldn't fit on the CD.

Updates: This repository contains newer versions of the packages in the Main Distribution repository. It's the server that's used whenever you attempt to update your system.

Backports: A Backports server allows access to software that's intended to go into the next version of Ubuntu but which has been packaged for the current version. This software might not have been tested thoroughly and so is suitable for neophiliacs or those who absolutely need the latest version (perhaps because of a vital new feature it offers).

Security Updates: Although the Updates server supplies newer versions of the main Ubuntu packages, sometimes security flaws need to be fixed almost immediately. If so, the packages concerned will be placed on this server. The Security Updates server isn't about new versions or functionality. It's about fixing security holes.

Note Sometimes, you might choose to add an additional repository. For example, in Chapter 28, you added the official Wine repository so you could install the latest version of the software that lets you run Windows programs.

Repository Sections

In addition to the categories listed in the previous sections, the Ubuntu repositories are split into *sections* according to how essential the software is to a basic Ubuntu installation, or the license that the software uses.

Here are the sections under which software is filed within a repository:

Main: This section contains nearly all the software that's featured in a basic Ubuntu installation. As such, it's all Free Software, and every package is supported by the Ubuntu project. That means that updates are frequently provided to fix security holes, or simply to keep up with latest releases.

Note *Free Software* refers to software that's licensed under the GNU Public License (GPL). It doesn't mean that the software is free of charge, although that's nearly always the case.

Restricted: Although Ubuntu is mostly Free Software, it must include some components that have license agreements that are not compatible with the goals of Free Software. That's what you'll find in this section. Many Ubuntu installations simply won't work fully without software from the Restricted section. The chief examples are hardware drivers that are released only in binary form (that is, closed source).

Universe: This section might be referred to as "the rest," because it contains the majority of Free Software available at the present time. Much of it is raided from the massive Debian software library, although the packages are sometimes tweaked to work correctly under Ubuntu before being made available. Unlike Main and Restricted, the Universe section is not officially supported by the Ubuntu project, which means there's no guarantee that security flaws will be fixed. Nor is there any guarantee of updates.

Multiverse: As with the Restricted section, here you'll find software that's released under a software license incompatible with either the word or spirit of Free Software. However, unlike the software in the Restricted section, none of the software in Multiverse is considered essential to a default Ubuntu installation.

Note When you configured the Synaptic Package Manager back in Chapter 8, all the sections on all the repository servers were activated (with the exception of Backports). This means that you have access to the widest possible range of software, but it also means that some of that software might not be Free Software, or compatible with the goals of the Ubuntu project. In particular, some software in the Multiverse section is considered beyond the pale by some Linux enthusiasts. You should check any software in the Multiverse section to be sure that you're happy with its licensing terms.

Not all repositories contain all the sections listed here. For example, the Updates and Security Updates repositories contain only Main and Restricted sections, because the Ubuntu project doesn't update software in the Universe and Multiverse subcategories.

All of this might sound a little confusing at first. The best policy is to take a look at the Synaptic Package Manager's Software Preferences dialog box. Open the program (System ► Administration ► Synaptic Package Manager), click Settings ► Repositories on the menu, and then take a look at what appears. The headings in the list are the repositories that you can access. Any with a check alongside them are those that APT, and therefore the Synaptic Package Manager, use. Those without a check are ignored.

Underneath each repository heading are the sections on that server that you can access. For example, the first repository in the list on my test PC is CD Ubuntu 5.10 "Breezy Badger" (Binary). This refers to the repository contained on the installation CD. Beneath that are the sections Officially Supported (Main) and Restricted Copyright (Restricted).

Note At the time of writing, the Synaptic Package Manager isn't clever enough to probe a repository to find out what sections are available. As you'll see if you select an online repository within the Software Preferences dialog box and click the Edit button, each section must be manually entered into the Sections field.

The Synaptic Package Manager in Depth

The Synaptic Package Manager is effectively a graphical front-end for the APT system. Later on in this chapter, we'll look at the command-line APT tools. Here, we'll take a closer look at using the Synaptic Package Manager.

Searching for Software

Before searching for software, it's nearly always a good idea to refresh the package databases. Just click the Reload button on the Synaptic Package Manager toolbar to grab the latest package lists from the various repositories that are in your sources.list file. Reloading can take a few

minutes on a slow connection, but it ensures that you have access to the latest software within the repositories.

To search for software, you can click any entry in the list of packages and simply start typing. This will match what you type against the package names in the list. Alternatively, you can click the Search button in the toolbar. By default, this searches through both package names and the descriptions, for a higher chance of a match. You can either type the specific program name or a keyword that may be within the description. For example, if you were looking for graphics drivers for your ATI Radeon card, but you didn't know the name of the package that contains them, you could type `radeon`.

Tip You don't need to type whole words in the search field. You can type part of a word or, more commonly, the word in a shortened or alternate form. For example, if you're looking for an e-mail client, it might be more fruitful to simply type "mail client," or even just "mail." This will then return results containing *e-mail*, *mail*, *mailing*, and so on.

By clicking the Settings ► Filters button, you can enhance your search by filtering out any packages that don't meet your requirements. It's advisable to click the New button to create your own filter before starting, as shown in Figure 29-3, rather than editing one that's already there. To apply a filter to search results, click the Custom button at the bottom left and then click the name of your filter in the list.

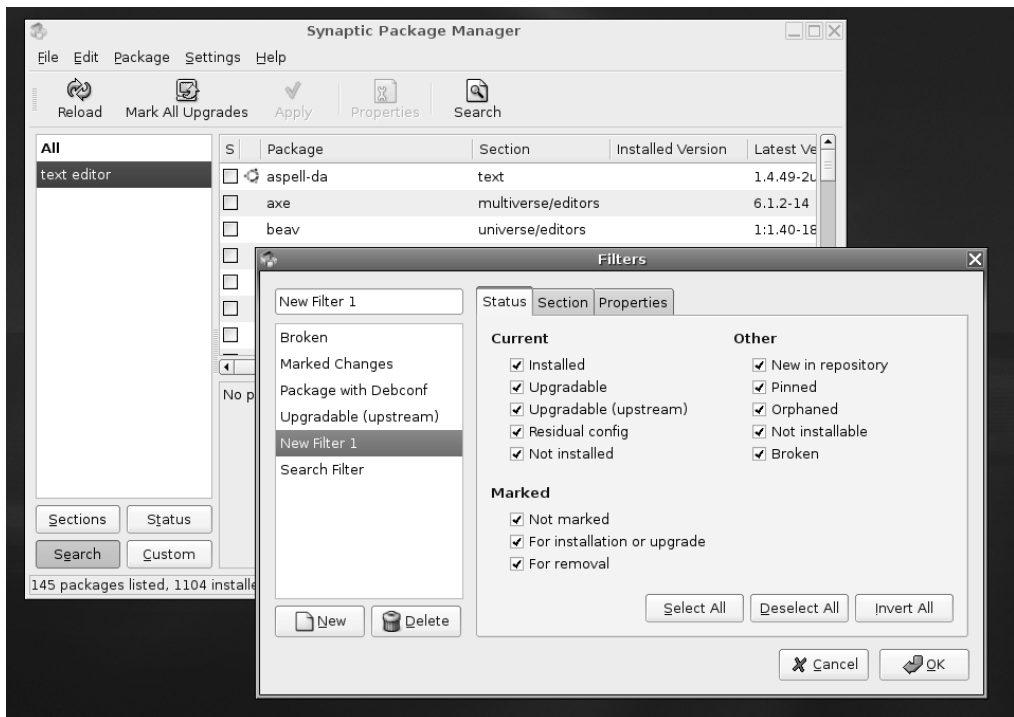


Figure 29-3. Filters can be used to trim the list of search results according to certain criteria.

One use of filtering is to remove the check alongside Installed so that you can remove from the search list any packages that might be already on your system. By clicking the Section tab, you can opt to search for only certain types of programs, such as those that are categorized as graphics programs or those that are categorized as communications packages (hold down the Ctrl key to search more than one package at once).

Note Filtering can help reduce the number of search results if you use a generic search term, but don't forget to deactivate it when you've finished. Click All at the top of the filters list, or simply click the Search button at the bottom left, to deactivate the filters.

Installing Software

When you click the check box next to a piece of software in the search results and select Mark for Installation, you'll be informed of any uninstalled dependencies that the program needs and asked to confirm their addition to the list of packages to be installed. If you right-click the file and selecting Mark Suggested for Installation or Mark Recommended for Installation, you'll see a list of programs that, although not essential to the running of the program in question, will enhance its features to some degree.

For example, if you choose to install the mutt e-mail program, it's also recommended that you install locales and mime-support, so Mutt will have multiple language support and will be able to better handle file attachments. You don't have to install these recommended programs; the software will run fine without them. But it can often be rewarding if you do so.

Note If the software in the recommended and suggested lists is grayed out, that means it's already installed. This happens frequently because recommended or suggested software is often required by other packages.

After you've made your selection of software (bear in mind that you can install more than one piece of software at once), click the Apply button on the toolbar. You'll see the Summary dialog box, as shown in Figure 29-4.

Here, you're once again asked to confirm what needs to be installed. If any software needs to be removed in order to meet dependency issues, you'll be told about this, too. Additionally, you'll be shown the total size of the files that will be downloaded, as well as the anticipated impact on your hard disk in terms of size after the programs are installed. At the very bottom of the Summary dialog box, you'll see a check box marked Download Package Files Only. As it suggests, this will download but not install the packages. If you then select the package for installation again in the future, you won't need to download it, and installation will be instantaneous (unless a newer version of the package has been released; in which case, the newer version will be downloaded and installed).

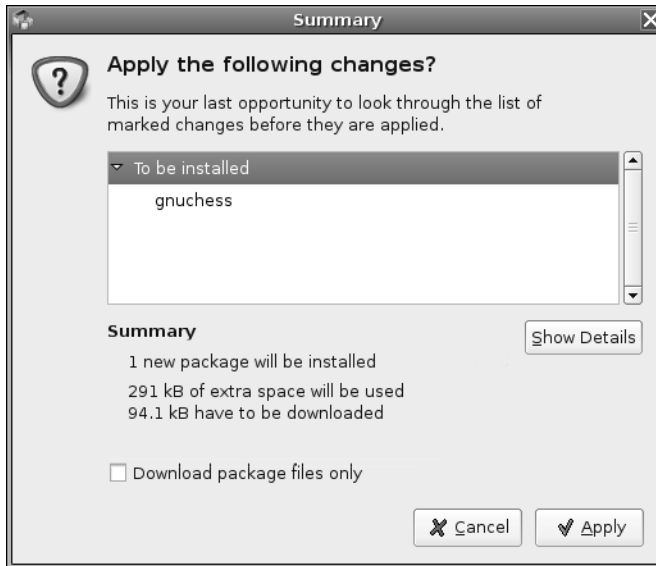


Figure 29-4. Before any software is installed by the Synaptic Package Manager, you'll be told what it is and asked to confirm the choice.

METAPACKAGES

Software such as the GNOME desktop actually consists of a number of programs and system libraries, rather than one single piece of software. Therefore, you might be wondering how, as just one example, you might install the KDE desktop under Ubuntu 5.10. Is it necessary to install each component's package manually?

In theory, dependency management should be able to help, and you should be able to select one key part of the KDE desktop system, such as the Konqueror file browser, and have the Synaptic Package Manager take care of the rest. After all, Konqueror will be dependent on other KDE packages.

Alas, this rarely works in reality. Installing Konqueror in this way will indeed install much of the KDE desktop suite, but not everything. Konqueror isn't reliant on Kate, for example, which is the default text editor under KDE. Perhaps more important, although the packages will be installed, there's no guarantee that they'll be configured to work correctly as a desktop environment.

Metapackages provide the solution. These are packages that contain configuration files to ensure the full range of software is installed and configured correctly, and they also have extensive lists of dependencies that include the complete set of packages for the software in question. (For what it's worth, the metapackage for KDE is `kdebase`.)

Alongside desktop suites, other examples of Ubuntu metapackages include the OpenOffice.org office suite, where the metapackage ensures all the components of the suite can be easily installed, and the X.org graphical subsystem. To see what metapackages are available, simply search for `metapackage` using the Synaptic Package Manager.

Removing Software

To remove a piece of software, search for it by name, then click the check box alongside it, and then select Mark for Removal. This will remove the software but leave behind any configuration files it created. This means you can install it again in the future, and it will function as it did before removal. However, you can also select Mark for Complete Removal, which will remove the configuration files.

As with installing software, the Synaptic Package Manager will attempt to manage dependencies when you remove software, but in this case, it will enforce the removal of any software that explicitly relies on that software.

This isn't an issue most of the time, but unfortunately, some packages have major reverse dependencies. For example, if you decide that you've had enough of the Base database program and want to remove it, you'll find that you'll also need to remove the entire OpenOffice.org suite, as well as the `ubuntu-desktop` metapackage.

Sometimes, packages have what seem like weird dependencies. For example, if you try to remove the Firefox web browser, you'll see that you need to remove the seemingly unrelated `gnome-app-install` package, as well as `yelp`, the GNOME online help system. The reason for this is that both these packages use a component of Firefox in some way and simply can't function without it. There isn't a hard-and-fast solution for situations like this. However, in many cases, you might simply choose to not remove the software and leave it on your system.

Package Management from the Command Prompt

Synaptic is one of the best examples of package-management programs around, and there's little reason to shun it and choose to install packages from the command line. However, you may find occasions to use `dpkg` or the APT tools. For example, if you're already working at the command line, then this method is quicker than starting up the Synaptic Package Manager.

Using `dpkg`

The most basic package-manipulation command is `dpkg`. `dpkg` allows you to perform a lot of package-related tasks, such as build packages from scratch. Here, we'll look at just simple package installation, removal, and query functions.

Note `dpkg` requires superuser powers to install or remove software, so must be preceded with `sudo`. But it can be run without superuser powers if you simply wish to query the package database. The same is true of the APT tools discussed later.

Installing Packages

`dpkg` is useful when you've already downloaded a specific `.deb` package and would like to install it. Here is the command:

```
sudo dpkg -i packagename.i386.deb
```

You must specify the entire filename, rather than just the name of the program.

`dpkg` is quick and dirty, and although it will warn you about any dependency issues, it will still go ahead and install the package. After installation, it will run the package's configuration scripts. But if there are missing dependencies, it won't be able to configure the program to work on your system, because it probably won't be in a usable state, as shown in the example in Figure 29-5.

A terminal window titled "keir@ubuntu: ~/Desktop" showing the execution of the command `sudo dpkg -i abiword-gnome_2.4.1-1ubuntu1_i386.deb`. The output shows the package being selected and unpacked, but then several dependency warnings are displayed: `dpkg: dependency problems prevent configuration of abiword-gnome: abiword-gnome depends on libenchant1c2; however: Package libenchant1c2 is not installed. abiword-gnome depends on abiword-common (= 2.4.1-1ubuntu1); however: Package abiword-common is not installed.` The process ends with `dpkg: error processing abiword-gnome (--install): dependency problems - leaving unconfigured` and `Errors were encountered while processing: abiword-gnome`. The prompt returns to `keir@ubuntu:~/Desktop$`.

```
keir@ubuntu: ~/Desktop
File Edit View Terminal Tabs Help
keir@ubuntu:~/Desktop$ sudo dpkg -i abiword-gnome_2.4.1-1ubuntu1_i386.deb
Password:
Selecting previously deselected package abiword-gnome.
(Reading database ... 63720 files and directories currently installed.)
Unpacking abiword-gnome (from abiword-gnome_2.4.1-1ubuntu1_i386.deb) ...
dpkg: dependency problems prevent configuration of abiword-gnome:
 abiword-gnome depends on libenchant1c2; however:
  Package libenchant1c2 is not installed.
abiword-gnome depends on abiword-common (= 2.4.1-1ubuntu1); however:
  Package abiword-common is not installed.
dpkg: error processing abiword-gnome (--install):
 dependency problems - leaving unconfigured
Errors were encountered while processing:
 abiword-gnome
keir@ubuntu:~/Desktop$
```

Figure 29-5. `dpkg` will warn about missing dependencies but will still install the package.

If this situation arises, it's up to you to install the dependencies that `dpkg` lists. Of course, some of these will have their own dependencies, which will also need to be installed. This cascade situation is informally known as *dependency hell* and is the main reason why the APT system was invented. As you'll see in the next section, APT effortlessly handles dependency issues like this.

If the dependencies aren't met after a `dpkg` installation, whenever you run the Synaptic Package Manager or attempt to use the APT tools, you'll be told of "broken" packages or unmet dependencies. APT will refuse to install any other software until the problem is fixed.

Tip Within the Synaptic Package Manager, you can click the Custom button at the bottom left and then click the Broken entry in the filter list in order to see any packages that have unmet dependencies. To repair the damage, click the check box of the package in question and then click Mark of Reinstallation. At the console, you can type `sudo apt-get -f install`. This will install all of the missing dependencies on the system.

Uninstalling Packages

To remove a package, type the following:

```
sudo dpkg -r packagename
```

Note that you simply need to type the name of the program, without its version number or the `.i386.deb` file extensions.

In this case, `dpkg` is slightly better behaved than when installing software. If there are any reverse dependencies (other packages that depend on the one you're trying to remove), you'll be stopped in your tracks with a couple of error messages. You'll need to resolve the reverse dependencies first but, of course, they might also have their own reverse dependencies . . . Welcome back to dependency hell!

Note The `dpkg -r` command will remove the package but leave behind its configuration files. This is handy if you install the software again in future. In order to remove the configuration files in addition to the software, type `sudo dpkg -P packagename`.

Querying Packages

`dpkg` includes a couple of query facilities that display details about packages. You can find out if a package is installed by typing this:

```
dpkg -l packagename
```

If you want to find out every bit of information about an installed package, including what dependencies it has, use the following command:

```
dpkg -s packagename |less
```

This example pipes the output of `dpkg` into `less` so you can read it more easily, because it's likely to fill several terminal window screens.

You can also use `dpkg` to query an installation file you've just downloaded:

```
dpkg -I packagename.i386.deb |less
```

All said, `dpkg` is an often undervalued tool that's capable of some handy low-level package-management tasks. Take a look at its man page to learn more.

Using the APT Tools

`dpkg` is the only option if you want to install a package file you've just downloaded. However, if you wish to use software repositories, you'll need to use the APT tools. These still use `dpkg` in the background to install and remove packages, but they also feature intelligence to handle dependency management.

Note If while using `dpkg` or `APT` you get an error message along the lines of, “Can’t get a lock,” make sure that the Synaptic Package Manager program isn’t open. Only one piece of software can access the package database at any one time.

Installing and Removing Packages

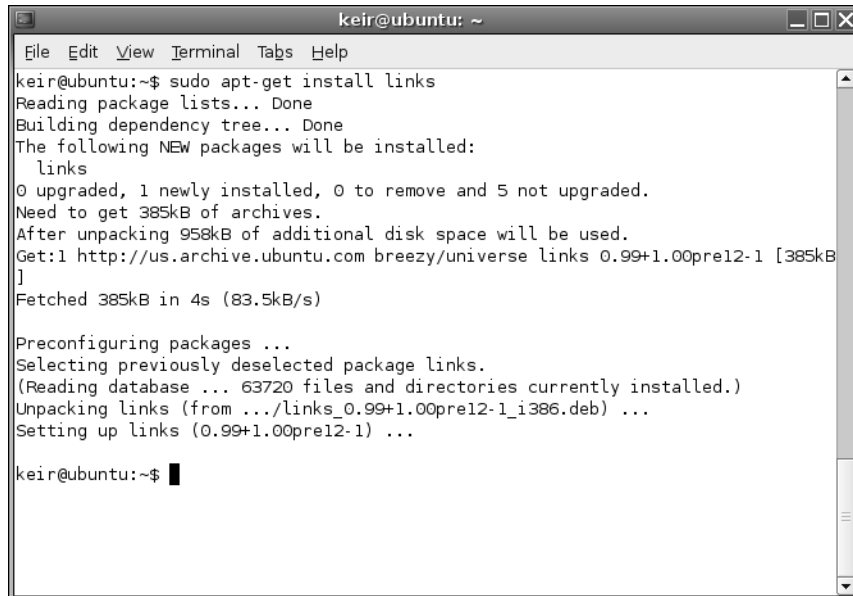
The most basic `APT` command is `apt-get`. You can use this command to install or remove packages contained within the repositories as follows:

```
sudo apt-get install packagename
sudo apt-get remove packagename
```

You should specify the program name without the version number. To install the `links` web browser, for example, you just need to type the following command:

```
sudo apt-get install links
```

Figure 29-6 shows the results. As you can see, `apt-get` will check dependencies, download the software, and then install it. It’s a much better way of working compared with `dpkg`.

A terminal window titled "keir@ubuntu: ~" showing the execution of the command "sudo apt-get install links". The output displays the process of reading package lists, building a dependency tree, and installing the "links" package. It shows that 385kB of archives are needed, and after unpacking, 958kB of additional disk space will be used. The package is fetched from a repository and then unpacked and set up on the system.

```
keir@ubuntu: ~
File Edit View Terminal Tabs Help
keir@ubuntu:~$ sudo apt-get install links
Reading package lists... Done
Building dependency tree... Done
The following NEW packages will be installed:
 links
0 upgraded, 1 newly installed, 0 to remove and 5 not upgraded.
Need to get 385kB of archives.
After unpacking 958kB of additional disk space will be used.
Get:1 http://us.archive.ubuntu.com breezy/universe links 0.99+1.00pre12-1 [385kB]
]
Fetched 385kB in 4s (83.5kB/s)

Preconfiguring packages ...
Selecting previously deselected package links.
(Reading database ... 63720 files and directories currently installed.)
Unpacking links (from .../links_0.99+1.00pre12-1_i386.deb) ...
Setting up links (0.99+1.00pre12-1) ...

keir@ubuntu:~$ █
```

Figure 29-6. You can use `apt-get` to install, remove, and update packages at the command line.

Note You can specify two or more programs to be installed and/or removed at the same time. Just separate the package names with a space: `apt-get install package1 package2 package3`, and so on.

It's a similar situation when it comes to uninstalling software. For example, suppose you tried to remove the Firefox web browser, like so:

```
sudo apt-get remove firefox
```

`apt-get` would also mark for removal `gnome-app-install` and `yelp`, two packages that depend on the browser. But before doing anything, it will tell you what it is about to do and ask you to confirm it.

Similarly, if you tried to install the AbiWord word processor, like so:

```
sudo apt-get install abiword-gnome
```

You would be informed that two additional packages need to be installed: `abiword-common` and `libenchant1c2`. They would be automatically added to the list of packages that were to be installed.

As with the Synaptic Package Manager, `apt-get` will also list suggested and recommended packages that will complement the software you wish to install but aren't vital. However, if you wish to install those packages, you'll need to do that later in a separate `apt-get` command.

Tip An alternative to the command-line APT tools is `aptitude`. This can be used like APT tools such as `apt-get`, but can also take into account suggested and recommended packages. For more information, see its man page.

Updating the System

`apt-get` can also perform various types of system updates. To update all the packages on your system to the latest versions contained within the repositories, type the following:

```
sudo apt-get upgrade
```

This is the command-line equivalent of using the Software Updates function of the Ubuntu desktop.

To upgrade the system to the latest version of Ubuntu, if there is one, type this:

```
sudo apt-get dist-upgrade
```

Perhaps it goes without saying that updating your system can take a long time, depending on the number and size of files to be downloaded. In addition, each package will need to configure itself during installation, and this can also take a long time.

Querying Packages and Repositories

To search the repository databases for particular software packages, use the `apt-cache` command:

```
apt-cache search packagename
```

Both descriptions and package names are searched.

You can also find out about dependencies using `apt-cache`:

```
apt-cache depends packagename |less
```

Once again, it's a good idea to pipe the output into `less` because the dependency list may run off the screen.

You can read the program description for a package file by typing the following:

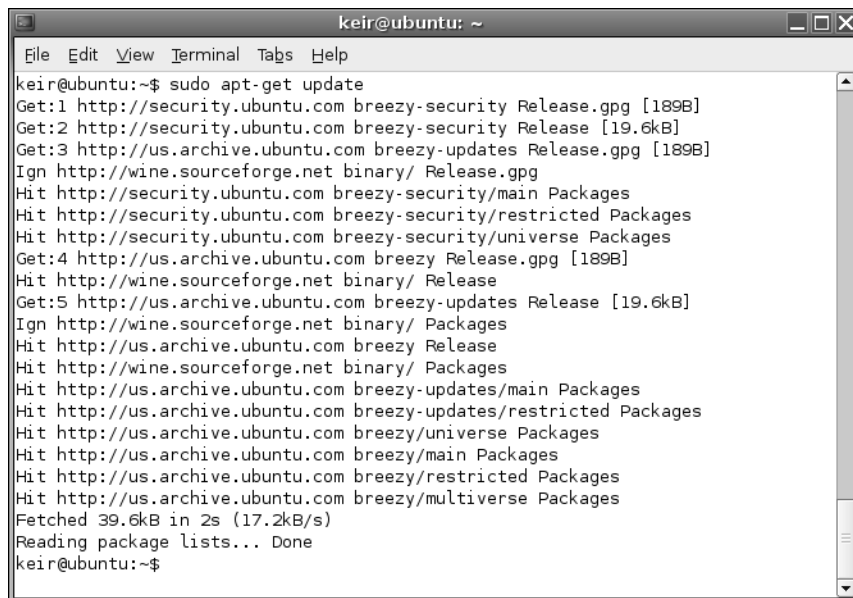
```
apt-cache show packagename |less
```

None of these commands makes a distinction between packages that are installed or otherwise—you're accessing the details held in the repository databases.

Before searching for packages, it's a good idea to make sure you have the latest package lists from the databases (the equivalent of clicking the Reload button in the Synaptic Package Manager). To refresh the lists, use this command:

```
sudo apt-get update
```

You should see output similar to that in Figure 29-7.



```
keir@ubuntu: ~
File Edit View Terminal Tabs Help
keir@ubuntu:~$ sudo apt-get update
Get:1 http://security.ubuntu.com breezy-security Release.gpg [189B]
Get:2 http://security.ubuntu.com breezy-security Release [19.6kB]
Get:3 http://us.archive.ubuntu.com breezy-updates Release.gpg [189B]
Ign http://wine.sourceforge.net binary/ Release.gpg
Hit http://security.ubuntu.com breezy-security/main Packages
Hit http://security.ubuntu.com breezy-security/restricted Packages
Hit http://security.ubuntu.com breezy-security/universe Packages
Get:4 http://us.archive.ubuntu.com breezy Release.gpg [189B]
Hit http://wine.sourceforge.net binary/ Release
Get:5 http://us.archive.ubuntu.com breezy-updates Release [19.6kB]
Ign http://wine.sourceforge.net binary/ Packages
Hit http://us.archive.ubuntu.com breezy Release
Hit http://wine.sourceforge.net binary/ Packages
Hit http://us.archive.ubuntu.com breezy-updates/main Packages
Hit http://us.archive.ubuntu.com breezy-updates/restricted Packages
Hit http://us.archive.ubuntu.com breezy/universe Packages
Hit http://us.archive.ubuntu.com breezy/main Packages
Hit http://us.archive.ubuntu.com breezy/restricted Packages
Hit http://us.archive.ubuntu.com breezy/multiverse Packages
Fetched 39.6kB in 2s (17.2kB/s)
Reading package lists... Done
keir@ubuntu:~$
```

Figure 29-7. Before searching for new packages, use the `apt-get update` command to ensure your package database is up-to-date.

DECODING PACKAGE FILENAMES

Although the filenames of packages might seem like cryptic mumbo-jumbo, they actually tell you a great deal about the file. Let's take a look at the package file of the Eye of GNOME image viewer to explain this:

```
eog_2.12.1-0ubuntu2_i386.deb
```

The first element of the filename is the name of the program. In this case, Eye of GNOME has been abbreviated to `eog`. Abbreviations like this are quite common because they decrease the length of the filename. But it's important to note that they will be consistent. For as long as Eye of GNOME is supported as a package under Ubuntu, its package filename will always begin with `eog`.

Following the name of the package is the version number of the program in question: `2.12.1-0`. This is the version number that will appear if you click Help ► About when the program is running, and is the version number decided on by the developer who created the software.

After the version number is the word `ubuntu`, which indicates that this is a package that's been created specifically for the Ubuntu distribution of Linux. Then you see the build version number of the package: `2`. This is Ubuntu's own version number, indicating how many times the package has been built (created) by the Ubuntu team. Sometimes, it's necessary to release an updated build of the same version of a program in order to correct an error that was accidentally introduced in the last build version. Sometimes, the program is patched by the Ubuntu team in order to support a new function.

After the version numbers is the platform on which the package will run. In this case, `i386` indicates that the package will run on all x86-based processors from the 80386 upwards (the 486, Pentium, Pentium II, AMD processors, and so on). Sometimes, you might see `i686`, which means that the package has been optimized for Pentium Pro chips and above (Pentium II, III, IV, and AMD's Athlon range of chips). If the package is created for 64-bit desktop processors, then `amd64` will appear there.

Optimized versions of packages for particular processors are used only when they might bring a performance boost. For example, there are `i686` versions of the Linux kernel and the `libc6` library. But it's rare for ordinary programs, like OpenOffice.org, to be optimized. This means the majority of packages that are used under Ubuntu have the `i386` designation.

Installing from Source

Back in the old days of Unix, the only way to install software was from source code, a process known as *compiling*. This was because most people edited the source code themselves, or at least liked to have the option of doing so. Nowadays, innovations such as the Debian package-management system make compiling all but redundant for the average user. But knowing how to compile a program from source is still a good Linux skill to have. In some cases, it's your only option for installing certain programs because you may not be able to find a packaged binary.

It goes without saying that program compilation is usually handled at the command prompt. It's not the kind of thing you would do via a GUI program.

Installing the Compiler Tools

Before you can compile from source, you need to install several items of software: the `make` program, which oversees the process of creating a new program, and the GNU Compiler Collection (GCC), which does the hard work of turning the source code into a binary. In addition,

if the software relies on certain library files, you'll need to install developer (`devel`) versions of them, as well as the libraries themselves if they're not already installed. For example, if you're compiling a program to run under the GNOME desktop, you'll need development versions of the GTK+ libraries.

Under Ubuntu, it's possible to install all the program-compilation tools you need by installing the `build-essential` metapackage. You can use the Synaptic Package Manager or the following `apt-get` command at the command prompt:

```
sudo apt-get build-essential
```

On my test system, most of these files were read from the installation CD, rather than downloaded, so you'll either have to have the CD present or remove its entry in your list of repositories in order to force APT to get the packages online.

Unpacking the Tarball and Solving Dependencies

Let's take a look at installing a program from source. Dillo is a stripped-down web browser that's designed for speed and small file size. It's a fun little program that's good to have around in the event of your main browser developing a glitch that you can't fix. The Dillo home page is www.dillo.org, so head over there and choose to download the latest version of the source code.

Note Okay, you got me. If you use `apt-cache` or the Synaptic Package Manager to look through the repositories, you'll see that Dillo is available as a ready-to-install package. But Ubuntu's package repositories are so comprehensive that, frankly, I couldn't find anything to demonstrate program compilation that wasn't already in there!

The first thing to do is to unpack and uncompress the tarball (if you wish to learn more about the `tar` command, see Chapter 32):

```
tar jxf dillo-0.8.5.tar.bz2
```

Of course, you should replace the filename with that of the version you downloaded.

Next, you'll need to switch into the source code directory and take a look at the `README` file. This will tell you what dependencies Dillo has and also any caveats you may need to take into account in order to compile Dillo on a Linux system:

```
cd dillo-0.8.5
less README
```

Note Unlike binary files, source code is rarely designed with one specific computer platform in mind. For example, Dillo is able to compile on all types of Unix, including Linux, Solaris, BSD, and others. With a little work, it might even be possible to compile it under Windows!

First, you see that Dillo needs the `glib` libraries. This is a given on nearly all Linux systems, but in order to compile, Dillo will need the `devel` version of `glib`, which isn't part of the default Ubuntu installation.

Next, you read that it also needs the GTK+ 1.2 libraries. These are present on the majority of GNOME-based Linux desktop systems but, once again, the `devel` versions will need to be installed.

Beneath that in Dillo's list of requirements is support for JPEG and PNG image formats, which are definitely installed on the average Linux system, and the `WGET` download tool, which is also included with most versions of Linux (although it's a good idea to use the Synaptic Package Manager or `apt-cache search` to check that it's installed).

After finding out about dependencies, you should scroll down the `README` to look for any notes about compiling under Linux. It turns out there might be some issues with older 2.4 versions of the Linux kernel, but Ubuntu uses 2.6, so this isn't an issue.

So, in short, before you can compile Dillo, you need to install `devel` versions of the `glib` and GTK+ 1.2 libraries. You can install these via the Synaptic Package Manager or `apt-get`. It will help cut down the search results if you realize that system library packages under Ubuntu are usually prefaced with `lib`. So, search for the `devel` versions of `libgtk` and `libglib`. Doing so on my test system returned three likely packages: `libglib1.2-dev`, `libglib2.0-dev`, and `libgtk1.2-dev`. There are two `libglib` entries because my system has both `glib2` and the older `glib1.2`.

To ensure compatibility, I decided to install `devel` versions of both. Since you're working at the command prompt, install the packages via `apt-get`:

```
sudo apt-get install libglib1.2-dev libglib2.0-dev libgtk1.2-dev
```

As soon as I typed this, it turned out that `libgtk1.2-dev` came with a host of dependencies in the form of X server `devel` libraries. The reasoning is that if the GTK+ `devel` library files are needed, these other libraries are often needed, too. Whatever the case, there's no harm in installing them.

Compiling

Now comes the exciting process of compiling the program! This is done via three commands, issued in sequence:

```
./configure  
make  
sudo make install
```

The first command starts the `configure` script, created by the Dillo programmer, which checks your system to ensure that it meets Dillo's requirements. In other words, it checks to make sure the `glib` and `GTK+` libraries are present. It also checks to make sure you have the correct software that's required to actually compile a program, such as `GCC` and `make`.

It's when the `configure` script is running that something is most likely to go wrong. In that case, more often than not, the error message will tell you that you're missing a dependency, which you must then resolve.

Note Some `configure` scripts are very thorough and check for components that the program you're trying to install might not even need. Because of this, you shouldn't worry if, as the text scrolls past, you see that various components are missing. Unless `configure` complains about it, it's not a problem.

The next command, `make`, takes care of the actual program compilation. When it's run, the screen will fill with what might look like gibberish, but this is merely the output of the GNU compiler. It provides a lot of valuable information to those who know about such things, but you can largely ignore it. However, you should keep your eyes peeled for any error messages. It's possible that the `configure` script might not have checked your system thoroughly enough, and you might be missing an important system component; in which case, `make` will halt.

Alternatively, the program simply might not be able to compile on your system without some tweaking to the `makefile` (the file that `make` uses). If such a situation arises, the best plan is to visit the web site of the developer of the software and see if there's a forum you can post to. Alternatively, check if the developer has an e-mail address you can contact to ask for help.

Eventually, the compilation will stop with a number of exit messages. Then the final command must be run: `make install`. This needs to be run with superuser powers because its job is to copy the binary files you've just created to the relevant system directories. In addition, any documentation that comes with the program is also copied to the relevant location on your system.

Once the three commands have completed, you should be able to run the program by typing its name at the command prompt. If you've been playing along at home and have compiled Dillo, you can run it by typing `dillo`, as you can see in Figure 29-8.

Note Perhaps it goes without saying that you'll have to add your own icon to the desktop or Applications menu. Source packages are usually designed to be installed on any version of Unix running a variety of desktop managers. It's therefore impossible for the developer to know where to create desktop shortcuts.



Figure 29-8. Dillo in action—a certain satisfaction comes from compiling a program from source code.

Summary

This chapter described how to install software under Ubuntu. We’ve looked at how this differs from Windows software installation, and how the Debian package-management system is designed to make life easier.

You learned how to use the Synaptic Package Manager to install software under the GUI, and how to use the `dpkg` and `APT` tools to install software at the command-line prompt. Finally, we looked at how programs can be compiled from their source code, which is a fundamental process of all versions of Linux.

In the next chapter, we’ll look at how to administer the system of users under Ubuntu.



Managing Users

Linux was designed from the ground up to be a multiuser system. When it is deployed on huge mainframe computers, it's capable of serving hundreds, if not thousands, of users at the same time, provided there are enough terminal computers for them to log in. In a more domestic setting, such as when Ubuntu is installed on a desktop PC, it usually means that more than one family member can have their very own account on the PC. Any files users create will be private, and users will also get their own desktop environment that is separate from that of the other users.

And even if you're the only person using your PC, you can still take advantage of Ubuntu's multiuser capabilities. Consider creating user accounts for various aspects of your life—perhaps one for work and one for time spent browsing the Web. Each user account can be tailored to a specific need.

In this chapter, you'll learn how to administer multiple user accounts.

Understanding User and Group Accounts

The concept of users and file ownership was explained in Chapter 14, but let's take a moment to recap and elaborate on some important points.

Users and Groups

Each person who wishes to use Ubuntu must have a user account. This will define what that user can and cannot do on the system, with specific reference to files. Because Ubuntu is effectively one large file system, with even hardware devices seen as individual files (see Chapter 14), this means that user permissions lie at the heart of controlling the entire system. They can limit which user has access to which hardware and software, and therefore control access to various PC functions.

Each user also belongs to a group. Groups have the same style of permissions as individual users. File access can be denied or granted to a user, depending on that person's group membership.

Note As in real life, a group can have many members and can be based around various interests. In a business environment, this might mean that a group is created for members of the accounting department, for example, or for the human resources department. By changing the permissions on files created by the group members, each group can have files that only the group members can access (although, as always, anyone with superuser powers can access all files).

On a default Ubuntu system with just a handful of users, the group concept might seem somewhat redundant. However, the concept of groups is fundamental to the way Ubuntu works and cannot be avoided. Even if you don't make use of groups, Ubuntu still requires your user account to be part of one.

In addition to actual human users, the Ubuntu system has its own set of user and group accounts. Various programs that access hardware resources or particular sets of files normally use these. Setting up system users and groups in this way makes the system more secure and easier to administer.

Root User

On most Linux systems, the root user has power over the entire system. Root can examine any file and configure any piece of hardware. Root typically belongs to its own unique group, also called root.

Ubuntu is different in that the root account is deactivated by default. Instead, certain users, including the one set up during installation, can “borrow” root-like powers by simply typing their login password. This is done by preceding commands with `sudo` at the command-line prompt, or as needed when using GUI programs that affect system settings.

If you wish, you can activate the root user account on your system and then log in to it when necessary. To activate the root account, use the following command:

```
sudo passwd root
```

After typing your own login password, you'll be invited to define a password for the root user.

Because of its power, the root user can cause a lot of accidental damage, so it's rare for anyone to log in as root on bootup. Instead, you can switch to root user temporarily from an ordinary user account by typing the following:

```
su
```

This will prompt you for the root password and then log you in as root for as long as you need. When you've finished, type `exit`, and you'll be returned to your ordinary user account.

Tip You can tell when you're logged in as root user because the command prompt will end with a hash (#). When logged in as an ordinary user, it ends with a dollar sign (\$). The hash symbol should be seen as a warning that you now have unrestricted control over the system, so be careful what you type and double-check everything before hitting Enter!

As an alternative to setting the root password, you can simply type the following whenever you want to switch to the root user account:

```
sudo su
```

You'll be prompted for your login password, in exactly the same way as if you had just preceded a command with `sudo`. After this, you'll be logged in as the root user. To quit the root user account, type `exit`.

UIDs and GIDs

Although we talk of user and group names, these are only used for the end user's benefit. Ubuntu uses a numerical system to identify users and groups. These are referred to as user IDs (UIDs) and group IDs (GIDs), respectively.

For various reasons, under Ubuntu, all the GID and UID numbers under 1000 are reserved for the system to use. This means that the first non-root user created on a system during installation will probably be given a UID of 1000. In addition, any new groups created after installation are numbered from 1000. On my system, the default user of `keir` had a UID of 1000 and a GID of 1000. The second user I added was given a UID of 1001 and a GID of 1001.

Note UID and GID information isn't important during everyday use, and most commands used to administer users and file permissions understand the human-readable usernames. However, knowing UIDs and GIDs can prove useful when you're undertaking more complicated system administration.

Adding and Deleting Users and Groups

The easiest and quickest way to add a new user or group is to use the Users and Groups tool under the System ► Administration menu. Of course, you can also perform these tasks through the command line.

Adding and Deleting Users via the GUI

To add a new user, select System ► Administration ► Users and Groups and click Add User. You'll see the User Account Editor dialog box, as shown in Figure 30-1.

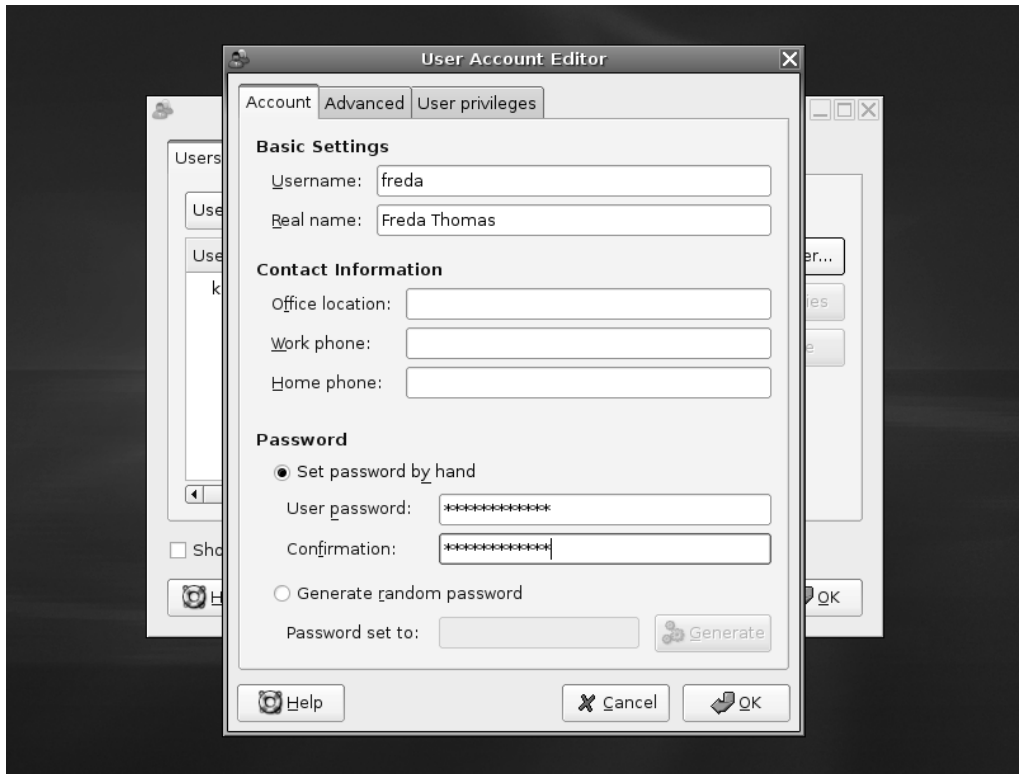


Figure 30-1. Adding new users and groups is easy courtesy of the Users and Groups program.

Fill out the fields on the Account tab, and optionally the Advanced and User Privileges tabs, as follows:

Account: As during initial installation, you're invited to enter a username for the user as well as the real name. The username is how the user is identified to the system, while the real name is how the user will be identified to other users. You can enter contact information under the relevant heading if you wish, but this isn't strictly necessary. However, you do need to set an initial password for the user. To ensure accuracy, enter it twice. Alternatively, the system can generate a random password from letters and numbers, but this may be harder for the user to remember.

Advanced: Here, you can alter various additional settings for the user. The chief one is assigning a user profile. The default for additional users is Desktop. This is a restricted account, which means that the user won't be able to administer the system or change settings. If you want the user to have these powers, click Administrator. Beneath this, you'll see various technical settings that can remain unchanged. However, you might like to change the main group for the user. By default, he will belong to a unique group based on his username (indicated by \$user). However, you could add the new user to your own group by selecting it from the list. Depending on your use of file permissions, this could grant the user access to your files.

Caution Many groups are listed in the Main Group drop-down list. Nearly all of these relate to the way the Linux operating system operates and can be ignored. You should never, ever delete any of these, or add a user to them!

User Privileges: The settings on this tab offer much more control over what a user can and cannot do on the system. Here, you can prevent users from using certain hardware, such as scanners or modems. You can also control whether the user is able to administer the system. Simply put a check alongside any relevant boxes.

Deleting a user is simply a matter of highlighting the username in the list and clicking the Delete button. Note that the user's `/home/` directory won't be deleted. You must do this manually with superuser powers, and it's best accomplished from the command-line prompt (see Chapter 14 for an introduction to basic file-manipulation shell commands).

Creating and Deleting Groups via the GUI

Adding a group is simply a matter of clicking the Groups tab in the Users and Groups program window (System ► Administration ► Users and Groups). After clicking the Add button, you'll be prompted to give the group a name. The group ID (GID) will be automatically filled in for you, but you could choose a different number if you have good reason to do so (remember to keep it above 1000 to keep in line with the way Ubuntu operates).

It isn't essential that you add users to the group there and then but, once again, you'll see a lot of user accounts in the list that belong to the system and are vital to the way Ubuntu operates. You shouldn't add any of these to your new group, and you should never, ever delete any of these user accounts!

As with user accounts, deleting a group is simply a matter of highlighting it in the list and clicking the Delete button. You should ensure that the group no longer has any members before doing this because, perhaps surprisingly, Ubuntu won't prevent you from deleting an active group (although it will warn you that this is a bad thing to do). Deleting an active group has the effect of preventing all users who belong to it from logging in.

Adding and Deleting Users and Groups at the Command Line

You can create new users at the command-line shell by using the `useradd` command. This command must be run with superuser powers, which is to say that it must be prefaced with the `sudo` command.

The command to add a user is normally used in the following way:

```
sudo useradd -m <username>
```

The `-m` command option tells the command to create a home directory for the user. Used on its own, `useradd` merely updates system files with the new user's details and nothing else. There are several other useful command options, which can be discovered by a quick browse of the command's man page.

Creating a new user this way will automatically add him to the `users` group. However, the Ubuntu way of working is to give each user his own group based on his username. Therefore, you will always need to create a new group for the user *before* you create the user account itself, using the `groupadd` command, as follows:

```
sudo groupadd <groupname>
```

Then you need to specify this group with the `-g` switch when creating a new user:

```
sudo useradd -m -g <username> <groupname>
```

For example, the following command creates a user called `raymond` and adds him to the group `raymond`:

```
sudo useradd -m -g raymond raymond
```

There's another more annoying issue relating to groups when you're creating a user account at the command line. Most users are members not only of their own group, but also of several system groups. These groups relate to various hardware and software functions. For example, membership of the `audio` group is required if the user wants to be able to use the sound card and hear audio. This is necessary because of the way Linux works. Therefore, you need to add new users to these groups if they're to make full use of the system. These groups are described as *supplementary groups*.

Use the `id` command to display user and group information. On my test system, typing the following:

```
id keir
```

revealed the following groups:

```
uid=1000(keir) gid=1000(keir)
groups=1000(keir),4(adm),20(dialout),24(cdrom),25(floppy),29(audio),30(dip),44
(video),36(plugdev),104(lpadmin),105(scanner),106(admin)
```

All those after my main group, `1000(keir)`, are supplementary groups. For a list of what they do, see Table 30-1.

Table 30-1. *System Groups Within Ubuntu*

Group	Definition
<code>adm</code>	Used for system logging
<code>dialout</code>	Required for use of serial port devices, such as older modems
<code>cdrom</code>	Allows user to access CD/DVD-ROM
<code>floppy</code>	Allows user to access floppy disk drive
<code>audio</code>	Enables sound output for user
<code>dip</code>	Required for use of dial-up modems

Table 30-1. *System Groups Within Ubuntu*

Group	Definition
video	Activates video acceleration for the user
plugdev	Allows user access to removable storage, such as card readers, digital cameras, and so on
lpadmin	Allows user to access the printer
scanner	Allows user to access the scanner
tape	Allows user to access tape storage devices (used for backup purposes)
fax	Allows user to access the fax facility
admin	Gives user system administration abilities (superuser powers)

As you might have guessed, to manually add a user under Ubuntu, you must not only create a group and then add the user to it, but you must also add that user to the required selection of supplementary groups. Some are mandatory, such as audio, while others are optional, depending on how much freedom you want to afford the new user.

You can add a user to supplementary groups by using the `-G` switch with `adduser`. Here's how to add a new user called `raymond` to the system so that he is able to make full use of the system (having first created a group called `raymond`):

```
sudo useradd -m -g raymond -G adm,dialout,cdrom,floppy,audio,dip,video,plugdev,
lpadmin,scanner,tape,fax,admin raymond
```

Additionally, creating a new user using `useradd` won't automatically apply a password to the account. Ubuntu can't work with passwordless accounts, so until one is applied, the new account will be locked. You can add a password using the `passwd` command, as discussed in the next section.

Deleting a user is mercifully simple compared to this! Use the `userdel` command, as follows:

```
sudo userdel <username>
```

This won't remove `raymond`'s `/home/` directory, however. That must be done manually. Similarly, to delete a group, use the `groupdel` command:

```
sudo groupdel <groupname>
```

Adding and Changing Passwords

On a default Ubuntu installation, ordinary users are able to change their passwords at the shell. The command for any user to change his password is simple:

```
passwd
```

The user will be asked to confirm his current password, and then to enter the new password twice, to confirm that it has been typed correctly.

Alternatively, by adopting superuser powers, a user can change the password of another account:

```
sudo passwd <username>
```

This is necessary just after you create a new user account because it isn't given a password automatically. For obvious security reasons Ubuntu won't allow blank passwords.

You can enter just about anything as a password, but you should bear in mind some common-sense rules. Ideally, passwords should be at least eight characters long and contain letters, numbers, and even punctuation symbols. You might also want to include both uppercase and lowercase letters.

A number of command options can be specified along with the `passwd` command when it used with superuser powers. For example, the `-l` option will lock the specified account so that it can't be accessed (the `-u` option will unlock it).

Tip You can temporarily switch into any user account by typing `su <username>`. When you've finished, simply type `exit` to return to your own account.

Summary

In this chapter, we looked at the principles behind user and group accounts under Ubuntu. We've examined how user and group accounts can be created, edited, and deleted using the GUI, as well as the command-line prompt. We also looked at how passwords can be manipulated by the individual users themselves and by a user with superuser powers.

In the next chapter, we'll look at how the system can be optimized. You'll also learn about several interesting and important system tools.



Optimizing Your System

One slight problem with Ubuntu (and all Linux distributions) is that they take a “one-size-fits-all” approach—the default installation attempts to provide services for the every kind of user. While this provides the widest range of compatibility, it doesn’t always ensure an optimized system.

You may never attach a printer to Ubuntu, for example, so what’s the point of keeping the printing subsystem in memory? You can remove it from your Linux setup and not only free memory, but also speed up boot times, because you no longer need to wait for the printer service to start. While this might save only a couple seconds, or just a couple of hundred kilobytes of memory, repeating the process and paring Ubuntu down to the bone can produce an ultra-efficient system.

In this chapter, you’ll learn how to target the various subsystems of your Linux system in order to optimize and speed up your computer. We’ll look at everything from bootup, to hard disks, to streamlining the kernel itself.

Speeding Up Booting

Let’s take a look at what happens when a Ubuntu-equipped PC boots. Then we’ll explore some ways to speed up the process.

Understanding Bootup

When you start your computer, initially, the computer’s BIOS searches for a boot program on the hard disk. In the case of Ubuntu, the boot program runs the GRUB boot loader. If you’ve installed Ubuntu alongside Windows (or any other operating system), the GRUB menu will appear at this stage, and you’ll be able to choose which operating system to load. If only Ubuntu is on the hard disk, you’ll see a brief prompt for three seconds telling you that the GRUB menu will appear if you press a key.

However your system is set up, GRUB has the same fundamental function: it’s designed to load the Linux kernel. The kernel then starts the very first program that’s run on any Linux system: `init`. The principal job of `init` is to run a variety of *run-level scripts*, which load the hardware and software necessary for the full and correct functioning of the system.

Two sets of run-level scripts run at this time: system initialization scripts, which are contained in `/etc/rcS.d`, and numbered run-level scripts, which are found in `/etc/rcX.d` (where `X` is the number of the current run level).

Note Actually, the `/etc/rc` directories don't contain the scripts. They merely contain symbolic links to scripts, which are contained in `/etc/init.d`.

The initialization scripts take care of the basics of the system, ensuring that vital hardware and software services are started. Initialization scripts are considered critical in order for the system to run correctly.

Numbered run-level scripts are more optional. They provide services the user may or may not need, depending on how the computer will be used. For example, a numbered run-level script might start the printing service. Another numbered run-level script might start the GUI components. Some users may not need either of these, so they could be removed.

Ubuntu has seven groups of numbered run-level scripts, ranging from 0 through 6. Each defines the mode in which the computer is running. For example, run level 1 is single-user mode. This means that only one user can log in, and networking is disabled (usually, many other nonessential services are not activated either).

Note Run level 6 is reboot mode and exists simply to reboot the system, while run level 0 is halt mode and will shut down the system. You'll probably never come into direct contact with these run levels. Instead the programs you use to shut down or restart the system, such as the System ► Log Out option within GNOME, use these run levels.

On Ubuntu, run levels 2 through 5 are defined as *multiuser*. Technically speaking, this means that they allow more than one user to log on, but actually, they're the day-to-day running modes of the computer. Run level 2 is the default run level under Ubuntu and, just as the system initialization scripts are contained in `/etc/rcS.d`, run level 2 scripts are contained in `/etc/rc2.d`, as shown in Figure 31-1.

Note In fact, Ubuntu's run levels 2 through 5 are identical. Run levels 3 through 5 might be described as spares, existing merely for further expansion possibilities. For what it's worth, it's theoretically possible to utilize run levels 7, 8, and 9, but few people do so because 2 through 5 offer more than most users need.

You might think that once the run-level scripts have completed, the system is ready to be used. But that's not the case. Although you'll be able to log on when the run-level scripts have finished, the GNOME desktop has yet to start, and this, too, has its own set of initialization processes. It needs to start its own set of programs, such as notification area applets, which provide handy functions like on-screen volume control.

```

keir@ubuntu: /etc/rc2.d
File Edit View Terminal Tabs Help
keir@ubuntu:/etc/rc2.d$ ls
S05vbesave  S13gdm      S20hotkey-setup  S25bluez-utils  S90usplash
S10acpid    S14ppp      S20makeudev     S25mdadm        S99acpi-support
S10sysklogd S19cupsys  S20pcmcia       S89anacron      S99fetchmail
S11klogd    S19hplip   S20powernowd    S89atd          S99rminologin
S12dbus     S20apmd    S20rsync        S89cron         S99stop-bootlogd
keir@ubuntu:/etc/rc2.d$

```

Figure 31-1. The scripts for each run level are contained in the `/etc/rcX.d` directories, where `X` is the run-level number.

Once all that has finished, you can use the computer!

Because so much must take place for your system to come to life, booting Ubuntu can take some time. On my test system, it averaged between one and two minutes. Certainly, you can shave some time from this.

Reducing the Boot Menu Delay

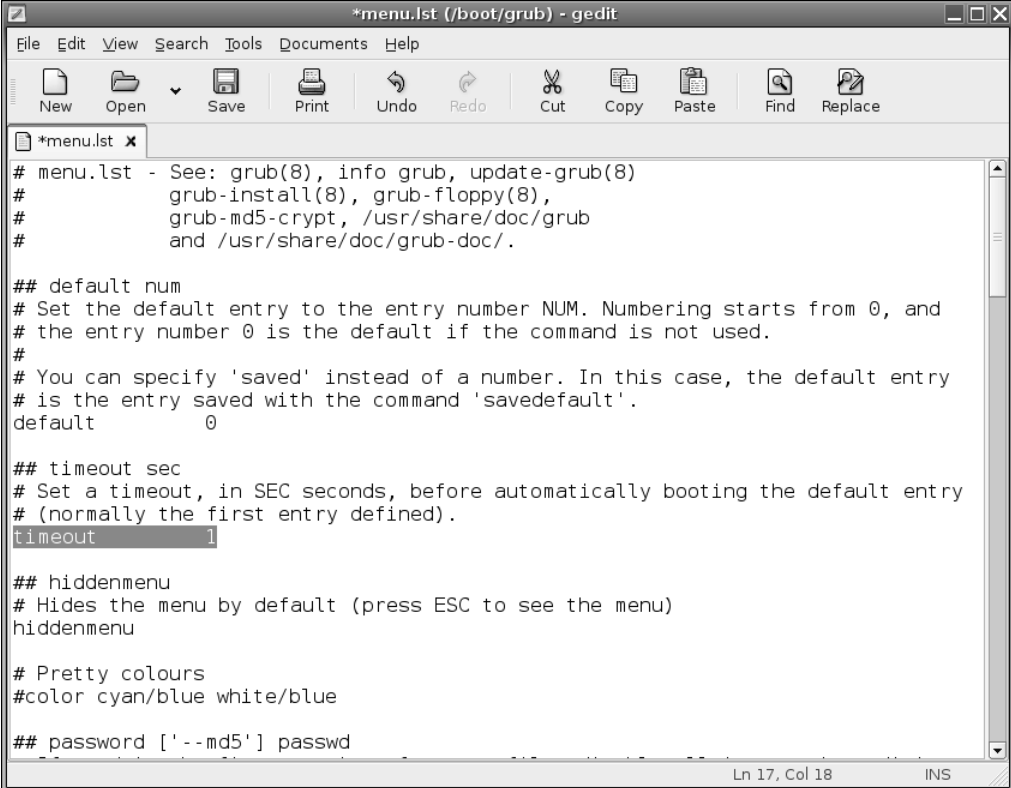
Getting rid of the GRUB boot menu delay can save some waiting around in the early stages of the boot process. The delay can be reduced to a second, or even eradicated completely. Of course, in such a case, you won't be able to choose which operating system you want to load if you're dual-booting with Windows. Even if Ubuntu is the only operating system on your computer, without the boot delay, you won't have the chance to boot into recovery mode, as offered on the GRUB menu. So you need to consider whether this is a worthwhile time-saving measure.

The boot menu delay is stated in the `/boot/grub/menu.lst` file. You can load this into the Gedit text editor by typing the following:

```
sudo gedit /boot/grub/menu.lst
```

Look for the line that begins with `timeout`, as shown in Figure 31-2, and change the value to whatever you wish. The units are counted in seconds, so a value of 3 equates to three seconds. A value of zero (0) will mean the boot menu won't appear at all. Generally speaking, a delay of a second (1) gives you just enough time to hit a key at the appropriate time, and this will then cancel the countdown, meaning the boot menu will stay on your screen until you select an option.

When you've finished, save the file and quit Gedit.



```
*menu.lst (/boot/grub) - gedit
File Edit View Search Tools Documents Help
New Open Save Print Undo Redo Cut Copy Paste Find Replace
# menu.lst - See: grub(8), info grub, update-grub(8)
# grub-install(8), grub-floppy(8),
# grub-md5-crypt, /usr/share/doc/grub
# and /usr/share/doc/grub-doc/.
## default num
# Set the default entry to the entry number NUM. Numbering starts from 0, and
# the entry number 0 is the default if the command is not used.
#
# You can specify 'saved' instead of a number. In this case, the default entry
# is the entry saved with the command 'savedefault'.
default 0
## timeout sec
# Set a timeout, in SEC seconds, before automatically booting the default entry
# (normally the first entry defined).
timeout 1
## hiddenmenu
# Hides the menu by default (press ESC to see the menu)
hiddenmenu
# Pretty colours
#color cyan/blue white/blue
## password ['--md5'] passwd
Ln 17, Col 18 INS
```

Figure 31-2. You can stop the GRUB menu hanging around for so long by changing the timeout value in its configuration file.

Optimizing Initialization and Run-Level Services

Perhaps it goes without saying that the majority of bootup time is spent starting the run-level scripts. This is when the entire system comes to life—hardware and essential software services are activated. But this isn't to say that all run-level scripts are essential.

Note A *service* is a piece of background software that provides something that you, the user, need on a day-to-day basis. Some services manage hardware, such as the graphical interface, printing services, and networking. Some services provide software services, such as logging files or checking the system clock against a time server.

The one-size-fits-all approach of Ubuntu means that some services that are started up aren't always necessary. A good example is the Bluetooth service. This is started up on every single Ubuntu system, yet only a fraction of users will ever use it. Therefore, if you don't use Bluetooth hardware (and are certain you never will), you can safely disable it and retrieve the chunk of memory it uses, as well as the amount of time it takes to start during bootup.

Approximately 60 run-level scripts start on a typical boot. By selective pruning, you can easily remove around a quarter or even a third of these, but caution is advised. You're altering a fundamental aspect of your system configuration, and one simple mistake can make the difference between a system that works and one that is no longer able to boot.

Disabling Run-Level Scripts

You can use the Services Settings program (System ► Administration ► Services) to control which run-level scripts start at bootup, but, sadly, it doesn't allow the enabling/disabling of initialization (run level S) scripts. It allows you to edit only certain numbered run-level scripts. Therefore, you need to download a command-line program called SysV Runlevel Config that can do the job. It offers a pseudo-graphical interface by which services can be activated and deactivated on all run levels, including S.

To obtain the SysV Runlevel Config program, use the Synaptic Package Manager. (If you haven't already set up the Synaptic Package Manager to use online repositories, see Chapter 8.) Select System ► Administration ► Synaptic Package Manager, click Search, and search for `sysv-rc-conf`. Mark it for installation, and then click Apply.

Open a GNOME Terminal window (Applications ► Accessories ► Terminal), and then maximize it to the full size of the screen. Then type the following to start the SysV Runlevel Config program:

```
sudo sysv-rc-conf -s 2S
```

This command runs SysV Runlevel Config showing only run levels 2 and S, to remove potential confusion between run levels.

The program's interface, shown in Figure 31-3, is simple. On the left, you see a list of the various scripts that are contained in the `/etc/init.d` directory and are therefore available for use during bootup (both initialization scripts and numbered run-level scripts). Not all of them are used. Those that aren't used are there in case they will be needed in future, or are provided for legacy reasons so that some software will work correctly. Along the top of the program window are the run levels you're going to edit: 2 and S.

If the check box next to a service has an X in it, that script is run on that particular run level. You can change this by navigating to the check box with the cursor keys and pressing the spacebar. You can scroll through the list of services by moving the cursor down to the bottom of the screen, or by pressing Ctrl+N to scroll down a page. Ctrl+P will move you up a page.

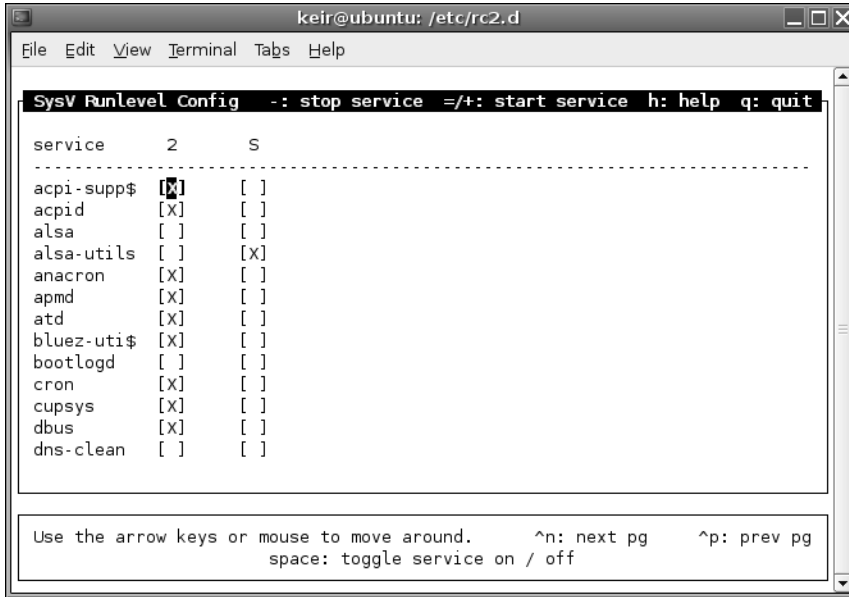


Figure 31-3. Remove the X from the check box alongside any service you would like to disable.

Table 31-1 lists the various run-level services, along with a short description of what they do and whether it's safe for you to disable them. *Only remove the check from the boxes of services you wish to deactivate!* Don't disable any services that you think should be running or otherwise interfere with the run-level settings. There's a very real risk you will damage your setup by doing so. Limit your changes to the services you specifically want to remove.

Note SysV Runlevel Config truncates the names of some services in the left hand column. In Table 31-1, I've listed the services as stated by the program, and then continued the full name within square brackets for clarity.

Table 31-1. Ubuntu Run Level Services

Service	Description
acpi-supply[ort]	Required for proper running of the Advanced Configuration and Power Interface (ACPI). It's vital to the correction functioning of Ubuntu, so don't disable it unless your computer uses the alternative Advanced Power Management (APM) system (see the <code>apmd</code> service). All modern computers use ACPI.
acpid	Also required for ACPI. Don't disable it.
alsa	At one time, this started the sound hardware, but it is no longer used. It's disabled by default in all run levels.
alsa-utils	Controls the volume of the system on startup. Don't disable it.

Table 31-1. *Ubuntu Run Level Services (Continued)*

Service	Description
anacron	Runs scheduled tasks, as described in Chapter 33. It allows the running of housekeeping tasks, so you shouldn't disable it (although it's not essential to the running of Ubuntu).
apmd	Activates the Advanced Power Management (APM) function. Nowadays, APM has been superseded by ACPI, so if you're sure your computer uses ACPI, you can disable this service.
atd	Scheduling service, as discussed in Chapter 33. If you don't use the <code>at</code> command, you can disable this.
bluez-uti[ls]	Required for the correct functioning of Bluetooth hardware (mobile phones, wireless headphones, and so on). If you don't use Bluetooth, you can disable this.
bootlogd	Logs boot messages. If you're sure you'll never, ever examine the boot log, as is sometimes necessary when problem solving, you can disable this service.
cron	Another scheduling service, as discussed in Chapter 33. Many housekeeping tasks are run via <code>cron</code> , so this should not be disabled (although it's not essential to the running of Ubuntu).
cupsys	Part of the printing subsystem. If you don't use a printer, either attached directly or across a network, this can be disabled.
dbus	Allows software to communicate with other software. It's vital, so don't disable it.
dns-clean	Part of the <code>ppp</code> subsystem that's designed to facilitate modem network connections. If you use a modem of any type, don't disable this. If you use an Ethernet or wireless connection, it can be disabled.
evms	Controls the Enterprise Volume Management Subsystem, which is used to manage multiple disk setups. If your computer doesn't use RAID or any other kind of logical partitioning scheme, you can disable this. If your computer has only one hard disk, you can definitely disable this service.
fetchmail	System service that can collect mail from remote mail servers for local delivery. It can be disabled on most systems, unless your mail client specifically uses it. The default e-mail client under Ubuntu, Evolution, doesn't use <code>fetchmail</code> .
gdm	Starts the GNOME desktop subsystem and therefore the X server. Don't disable this, unless you want a command-prompt-only run level.
halt	Used to shut down the system. Don't disable this service.
hdparm	Runs the <code>hdparm.conf</code> script, which allows you to optimize your hard disk settings. Don't disable this.
hotkey-se[utup]	Configures various additional keys on some notebook computers, such as those that allow the dimming of the screen. You can disable this, unless you're using a notebook computer.
hotplug	Controls systems such as USB and PCMCIA that allow hardware to be inserted when the computer is switched on. Also controls some sound devices. Don't disable this.

Table 31-1. *Ubuntu Run Level Services (Continued)*

Service	Description
hotplug-n[et]	Works like hotplug. Don't disable it.
hplip	Subsystem for Hewlett Packard (HP) scanners and/or printers. If you don't use an HP scanner or printer, you can disable this.
ifrename	Part of the networking subsystem that allows the renaming of network devices. Unless you're running a complicated server system, this can be safely disabled. If you find your network connection doesn't work afterwards, enable it.
ifupdown	Part of the network subsystem. Don't disable it.
ifupdown-[clean]	Part of the network subsystem. Don't disable it.
klogd	Logs kernel messages. This is handy for problem solving, but this service can be disabled if you're absolutely sure you'll never need it.
linux-res[tricted-modules-common]	Handles the loading of closed-source kernel modules, such as some wireless card and graphics card drivers. Unless you're sure you don't use such drivers, don't disable this service.
lvm	Controls the Logical Volume Management (LVM) subsystem. If your system doesn't use RAID or any other kind of LVM system, this service can be disabled.
makedev	Controls the /dev hardware subsystem. Don't disable it.
mdadm	Controls the software RAID setup. Disable it if you don't use software RAID.
mdadm-raid	Controls the software RAID setup. Disable it if you don't use software RAID.
module-in[it-tools]	Loads additional kernel module(s). Don't disable it.
mountvirt[fs]	Mounts virtual file systems that are essential to Ubuntu. Don't disable it.
networking	Manages network interfaces. Don't disable it.
ntpddate	Synchronizes time with the Ubuntu time servers. You can disable this, unless your PC's system clock is faulty and loses/gains time.
pcmcia	Starts PCMCIA devices on notebooks. You can disable this if you use a desktop computer, unless you have a PCI-based wireless card (or similar) that uses a PCMCIA card.
powernowd	Controls the CPU frequency scaling service that's common on many notebook and desktop computers. This can be disabled if absolutely necessary, but it's a good thing to have around.
ppp	Controls PPP connections, such as those established by modems (both dial-up and DSL/cable). If you don't use a modem, you can disable this.
pppd-dns	Controls PPP connections. You can disable it if you don't use a modem.
readahead	Optimization software designed to store commonly used libraries in memory. It's designed to speed up your system, but many people report it makes little difference, so it can be disabled.

Table 31-1. *Ubuntu Run Level Services (Continued)*

Service	Description
reboot	Responsible for rebooting the system. Don't adjust this setting.
rmnologin	Script that removes a system script that prevents nonroot users from logging in during startup. Don't disable it.
rsync	Controls the rsync service that allows the efficient transfer of files between computers. This can be disabled.
screen-cl[eanup]	Takes care of some bootstrap housework. Don't disable it.
sendsigs	Kills processes during shutdown and rebooting. Don't disable it.
single	Runs scripts specific to the correct function of run level 1 (single-user mode). Don't alter this setting.
stop-boot[logd]	Related to bootlogd and can be disabled if bootlogd is.
sudo	Security feature to ensure the grace period that sudo allows before it once again asks for your password. This can be disabled, although the security conscious may like to leave it active.
sysklogd	Starts the system logging service. As with klogd, it can be disabled if you do not wish to enable system logging.
udev	Part of the hardware subsystem. Don't disable it.
udev-mtab	Part of the hardware subsystem. Don't disable it
umountfs	Unmounts file systems at reboot or shutdown. Don't disable it.
urandom	Controls the random number generator. Don't disable it.
usplash	Part of the system that causes the pretty graphics to appear during boot. You can disable this.
vbesave	Part of the graphical subsystem. Don't disable this.
xorg-comm[on]	Starts the graphical subsystem. Don't disable this. If you wish to deactivate the GUI, disable the gdm service.

When you've finished using SysV Runlevel Config, type `Q` to quit the program. Then reboot to test your new settings.

Repairing Problems

If you find you've introduced a terrible problem, the solution is simply to switch to another run level. Run levels 3 through 5 are clones of 2 (or they were before you edited run level 2), so you can simply make one of those the default. Then, after rebooting into it, run SysV Runlevel Control again to experiment and reintroduce some excluded services into run level 2. Then you can switch back to run level 2 afterwards.

The default run level is set in the `/etc/inittab` file. You can load this into the nano command-line text editor, like so:

```
sudo nano /etc/inittab
```

Then look for the line that reads `id:2:initdefault:` and replace it with `id:3:initdefault:` (to switch to run level 3; to revert to run level 2, replace 3 with 2). Press Ctrl+X to exit nano and answer Y to the option of saving the buffer, which will save the file.

Optimizing Hard Disk Settings

The hard disk is one of the key elements in the modern PC. Because most of your PC's data must travel to and from it, speeding up your hard disk means that your entire PC will be faster.

Ubuntu provides a powerful command-line tool that you can use to control every aspect of your hard disk: `hdparm`. This is a power-user's tool. Not only must it be run with superuser powers (preceded by `sudo` at the command line), but you also must be careful not to mistype the commands. All changes are made instantly, so if you make a mistake, your system may crash, or at least suffer from serious problems. There's even the risk of data loss, although this is minimized by making sure that you have no other programs running at the same time you run `hdparm`.

The good news is that changes made via `hdparm` will last for only the current session, so there's no risk of permanent damage. Any changes that are beneficial can be made permanent later.

In the context of optimization, `hdparm` lets you both benchmark the disk and change various technical settings, such as the sector `multcount` value. These adjustments can bring speed boosts.

Benchmarking Your Hard Disk

Because experimenting with `hdparm` can cause crashes, and because its benchmarking feature needs almost exclusive access to the hard disk, `hdparm` is best run with as few as possible additional programs up and running. Therefore, killing the GUI is a good idea. To stop the GUI, open a GNOME Terminal window (Applications ► Accessories ► Terminal), and type the following:

```
sudo /etc/init.d/gdm stop
```

You'll then need to log in again, so do so.

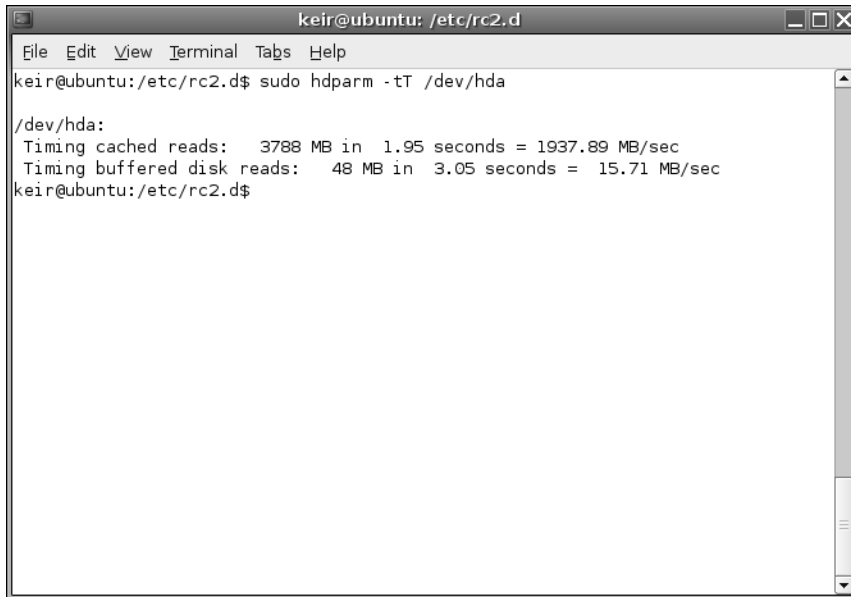
Note Technically speaking, switching to run level 1 is an even better idea, because this will deactivate all unnecessary services. Run level 1 is akin to the Windows Safe Mode, except without the GUI. However, you want realistic benchmark results to test the changes you make via `hdparm`, and it's debatable whether the restricted confines of run level 1 will provide such results.

Let's start by benchmarking your hard disk to see its performance based on the current settings. Type the following (assuming Ubuntu is installed on the first hard disk in your system; if it's on the second hard disk, change `/dev/hda` to `/dev/hdb`):

```
sudo hdparm -tT /dev/hda
```

This will benchmark your disk in two ways. The first tests the PC's memory throughput, measuring the data rate of the memory, CPU, and cache. The second actually tests the disk's

data rate. The second test affects the outcome of the first, which is why the two are used together. Between them, these two methods of benchmarking present the standard way your disk is used on a day-to-day basis. Figure 31-4 shows the results on my system.



```

keir@ubuntu: /etc/rc2.d
File Edit View Terminal Tabs Help
keir@ubuntu:/etc/rc2.d$ sudo hdparm -tT /dev/hda
/dev/hda:
Timing cached reads:  3788 MB in 1.95 seconds = 1937.89 MB/sec
Timing buffered disk reads:  48 MB in 3.05 seconds = 15.71 MB/sec
keir@ubuntu:/etc/rc2.d$

```

Figure 31-4. The `hdparm` program can be used to both benchmark and optimize your hard disk.

Make a note of the figures so that you can compare them to the results of these tests after you change hard disk settings.

Changing Hard Disk Settings

You can use `hdparm` to view your current hard disk settings by entering the following at the command prompt:

```
sudo hdparm /dev/hda
```

On my test PC, these are the results I got:

```

/dev/hda:
multcount      = 0 (off)
IO_support     = 1 (32-bit)
unmaskirq     = 1 (on)
using_dma     = 1 (on)
keepsettings  = 0 (off)
readonly      = 0 (off)
readahead     = 256 (on)
geometry      = 65535/16/63, sectors = 160086528, start = 0

```

Let's take a look at what these settings mean.

The Multcount Setting

The first, `multcount`, refers to how many sectors can be read from the hard disk at any one time. The theory is that the highest possible value here is best. Most modern hard drives support a value of up to 32. You can find out by issuing the following command:

```
sudo hdparm -i /dev/hda
```

Look for `MaxMultSect` in the results. On my test PC, this read `MaxMultSect=16`.

Ironically, although higher values are thought best, sometimes a lower value can speed up hard disk access. You can experiment with the `multcount` setting on your hard disk by using the `-m hdparm` command option:

```
sudo hdparm -m16 /dev/hda
```

You can then follow this by another benchmark to see if there is an improvement:

```
sudo hdparm -tT /dev/hda
```

If there isn't any improvement, you could try a lower value by simply swapping the `-m16` for `-m8`.

The IO_Support Setting

The `IO_support` line refers to the input/output (I/O) mode used by the hard disk controller. There are three possible settings: 0, to disable 32-bit support; 1 to enable 32-bit support; and 3 to enable 32-bit support with a special sync signal.

You can change the `IO_support` setting with the `-c hdparm` command option, and the 32-bit support with `sync` option (3) is generally considered the best choice:

```
sudo hdparm -c3 /dev/hda
```

The Unmaskirq Setting

The third setting, `umaskirq`, allows Ubuntu to attend to other tasks while waiting for your hard disk to return data. This won't affect hard disk performance very much, and generally it's a good idea for the health of your system to activate it if isn't already switched on. This command activates `umaskirq`:

```
sudo hdparm -u1 /dev/hda
```

The Using_Dma Setting

The fourth setting refers to whether Direct Memory Access (DMA) is in use. Hard disks are sold on the basis of their DMA modes, such as UltraDMA Burst 2 and the like. DMA is considered an indicator as to the speed of a hard disk, but the truth is that, like any specification, it is only a guide.

DMA is activated by default under Ubuntu, but you can alter the DMA mode using the `-X` command option. However, on most modern PCs, this isn't necessary because the computer's BIOS defaults to the fastest DMA mode.

Other Settings

The last three settings, above the summary of the geometry and sector information of the disk, are those you shouldn't change. The `readahead` setting controls how many hard disk blocks are loaded in advance. It doesn't affect the performance of modern IDE-based hard disks, because the drive electronics contain buffers that perform this task themselves.

The `keepsettings` setting refers to the ability of the drive to remember `hdparm` settings over a reboot, which isn't necessary because you intend to use the `hdparm.conf` file to change the disk settings at each boot. The `readonly` setting sets whether or not the hard disk is read-only (so that no data can be written to it). Changing this setting is not advisable!

Making Disk Optimizations Permanent

The `/etc/hdparm.conf` file is read at each bootup, and any settings it contains are applied to the hard disk. You can edit this file in the Gedit text editor by typing the following:

```
sudo gedit /etc/hdparm.conf
```

The simplest way of using this file is to edit the `hdparm` command string at the end of the file. To do so, start by finding the line that reads as follows:

```
#command_line {
```

Delete the hash mark from the beginning of the line. Then delete the next line (the one that starts with a hash mark, followed by `hdparm -q -m16 . . .`), and type your own version of the `hdparm` command, complete with the command options you discovered previously. For example, if your benchmarking has shown that the `multcount` and `IO_support` command options bring dividends, you might type something like the following:

```
hdparm -m16 -c3 /dev/hda
```

There's no need to type a hash mark at the beginning, as with the line you deleted. Then delete the hash mark from in front of the line beneath this, so that it's simply a bracket (`}`) on its own, as shown in Figure 31-5.

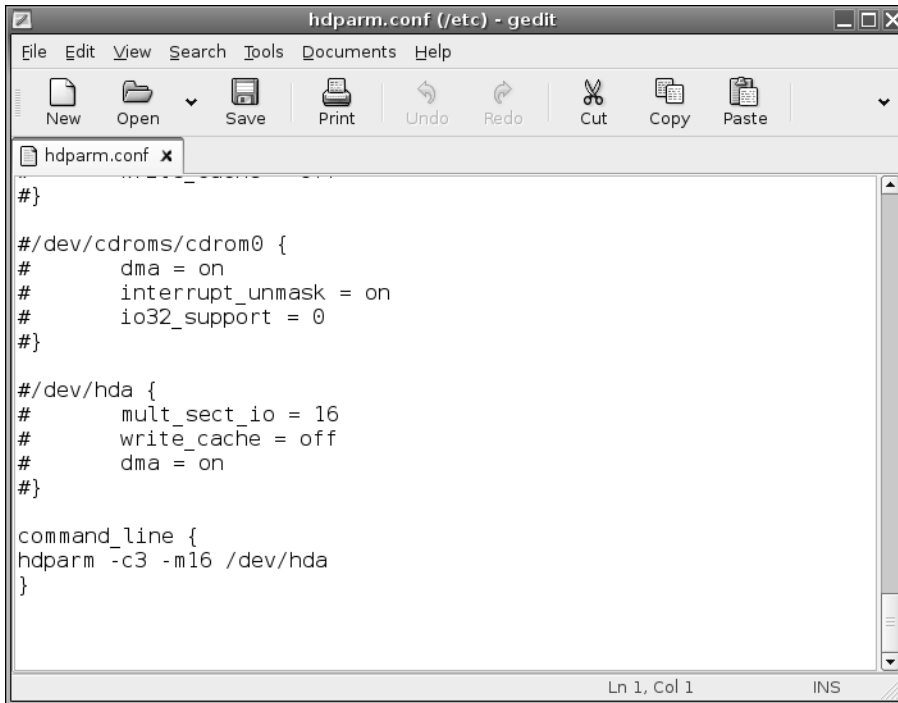


Figure 31-5. You can make your disk tweaking permanent by editing the `hdparm.conf` file.

When you've finished, save the file and reboot to test your settings.

Prelinking

As discussed in Chapter 29, a lot of Ubuntu software relies on other pieces of code to work. These are sometimes referred to as *libraries*, which is a good indicator of their purpose: to provide functions that programs can check in and out whenever they need them, as if they were borrowing books from a library.

Whenever a program starts, it must look for these other libraries and load them into memory so they're ready for use. This can take some time, particularly on larger and more complicated programs. Because of this, the concept of *prelinking* was invented. By a series of complicated tricks, the prelink program makes each bit of software you might run aware of the libraries it needs, so that memory can be better allocated.

Prelinking claims to boost program startup times by up to 50% or more, but the problem is that it's a *hack*—a programming trick designed to make your system work in a nonstandard way. Because of this, some programs are incompatible with prelinking. In fact, some might simply refuse to work unless prelinking is deactivated. At the time of writing, such programs are in the minority. However, keep in mind that prelinking can be easily reversed if necessary. Alternatively, you might want to weigh whether it's actually worth setting up prelinking in the first place.

Note Many of the GNOME programs under Ubuntu aren't compiled in a way that's compatible with prelinking. Therefore, you might not see much of a speed boost using the GNOME desktop.

Configuring Prelinking

If you decide to go ahead with prelinking, you'll need to download the relevant software from the Ubuntu software repositories. (If you haven't already set up the Synaptic Package Manager to use online repositories, see Chapter 8). Open Synaptic Package Manager (System ► Administration), click the Search button, and type `prelink` into the search box. Mark `prelink` for installation, and then click Apply.

Before you can run a prelinking sweep of your system, you need to enable it in one of its configuration files. To do this, type the following in a GNOME Terminal window:

```
sudo gedit /etc/default/prelink
```

Change the line that reads `PRELINKING=unknown` to `PRELINKING=yes`. Then save the file and quit Gedit.

To run a prelinking scan of your system, simply issue this command:

```
sudo prelink -a
```

This will prelink practically all the binary files on your system and may take some time to complete. You may also see some error output, but you don't need to pay attention to it.

Prelinking was automatically added as a daily cron job when you installed it (see Chapter 33 for a description of cron), so any new programs you add will be automatically prelinked.

Deactivating Prelinking

Should you find prelinking makes a particular application malfunction or simply stop working, you can try undoing prelinking. To do this, find out where the main binary for the program resides, and then issue the `prelink` command with the `--undo` command option. For example, to remove prelinking from the Gedit text editor program, you could type the following:

```
whereis gedit  
sudo prelink --undo /usr/bin/gedit
```

However, this may not work because some programs might rely on additional binaries on the system. Therefore, the solution might be to undo prelinking for the entire system, which you can do by typing the following:

```
sudo prelink -ua
```

After this, you should remove the `prelink` package to stop it running again in future (or manually remove its cron entry, as explained in Chapter 33).

OPTIMIZING THE KERNEL

You can download the Linux kernel source code and compile your own version of it. This gives you total control over the kernel configuration, so you can leave out parts you don't want in order to free memory. You can also set certain optimization settings, such as creating a version of the kernel specifically built for your model of CPU.

Although compiling a kernel is a simple procedure, there are many complex questions that you'll need to answer, and an in-depth knowledge of the way Linux works is necessary.

In addition, compiling your own kernel brings with several issues. The first is that it may not work with any binary modules that you have installed, such as graphics cards or wireless drivers. You can opt to install these yourself from scratch, but this adds to the complexity.

The second problem is that Ubuntu is built around precompiled kernels. Several software packages expect to work with the precompiled kernel and, in addition, Ubuntu may occasionally download an updated prepackaged kernel automatically as part of the system update feature and override the one you've created.

If there are any security problems with the kernel version you compiled, you'll need to recompile a new kernel from scratch (or patch the one you have). This means you'll have to keep an eye on the security news sites and take action when necessary.

That said, compiling a kernel is an excellent way of learning how Linux works, and the sense of achievement if it all goes well is enormous.

Some people choose to download the kernel source code from the official Linux kernel site, www.kernel.org. However, it makes more sense to download the official Ubuntu release, because this will be tailored for the way your system works. Using the Synaptic Package Manager, simply search for `kernel-tree`.

You can find several guides to compiling your own kernel online, but I recommend the following posting on the Ubuntu forums web site, which looks at compiling a kernel under Ubuntu: <http://ubuntuforums.org/showthread.php?t=85064>.

Freeing Disk Space

After using Ubuntu for some time, you might find that the disk begins to get full. You can keep an eye on disk usage by using the following command in a GNOME Terminal window:

```
df -h
```

This will show the free space in terms of megabytes or gigabytes, and also expressed as a percentage figure.

If the disk does start to get full, you can take some steps to make more space available.

Emptying the /tmp Folder

An easy way to regain disk space is to empty the `/tmp` folder. As with the Windows operating system, this is the folder in which temporary data is stored. Some applications clean up after themselves, but others don't, leaving behind many megabytes of detritus.

Because the `/tmp` folder is accessed practically every second the system is up and running, to empty it safely, it's necessary to switch to run level 1. This ensures few other programs are running and avoids the risk of deleting data that is in use. The following series of commands will switch to run level 1, empty the `/tmp` folder, and then reboot afterwards:

```
sudo killall gdm
[log in with your username and password]
sudo init 1
rm -rf /tmp/*
reboot
```

Tip On a similar theme, don't forget to empty the desktop Wastebasket. This can hold many megabytes of old data.

Emptying the Cache of Package Files

You might also choose to clear out the Advanced Packaging Tool (APT) cache of old .deb package files. On a system that has been very frequently updated, this can free many megabytes (possibly gigabytes) of space, as shown in Figure 31-6.




Figure 31-6. Over time, Ubuntu builds up a library of old package installation files, which you can delete to free some disk space.

You can empty the cache by typing the following command in a GNOME Terminal window:

```
sudo rm -f /var/cache/apt/archives/*.deb
```

Caution Be careful to type the command to empty the APT cache exactly as it's written. Even inserting an additional space can lead to very bad consequences!

If you want to restore any packages later on, simply locate them in the Synaptic Package Manager list, click the check box, and click Mark for Reinstallation. This will cause the package to be downloaded, installed, and configured.

Removing Unused Software

If you still need disk space, consider uninstalling unused programs. As you've learned, you manage software through the Synaptic Package Manager (System ► Administration).

To remove a package, click the check box and select Mark for Removal. However, it's not a good idea to simply scroll down the list and remove anything that seems dispensable. Because of the way Linux works, many seemingly insignificant packages are actually vital to the running of the system. Instead, it's a better idea to look for programs on the Applications menu, and then return to the Synaptic Package Manager to remove them by searching for their name.

As always, removing software can create dependency problems, so you might find yourself limited in what software you can actually remove.

Tip If you want to remove all the desktop games, simply search for `gnome-games` in the Synaptic Package Manager and mark it for removal.

Adding Another Disk Drive

Another solution to the problem of running out of disk space is to add another hard disk drive, perhaps as a slave on the primary IDE channel. Using a new disk within Ubuntu is very easy and can be done from the command-line shell.

Partitioning the Disk

Once the disk has been fitted, boot into Ubuntu and open a GNOME Terminal window. Use the `cfdisk` command to initially partition the disk. Assuming that you've added the new disk as a slave on the primary channel, use the following command:

```
sudo cfdisk /dev/hdb
```

Working out how Ubuntu refers to the hard disks installed on the system isn't hard. Usually, they're given letters from a through to d. So, `/dev/hda` is the primary master, `/dev/hdb` is the primary slave, `/dev/hdc` is the secondary master (usually the CD/DVD-ROM drive), and `/dev/hdd` is the secondary slave. If your system uses SCSI drives, you'll find they're named `/dev/sda`, `/dev/sdb`, and so on.

The first thing to do within the `cfdisk` program is create a new partition. Use the cursor keys to highlight New, and then press Enter. The default partition size should automatically be

all of the disk space, so press Enter again to confirm this. With the new partition created, highlight Write on the menu and press Enter. This will write the new partition information. After this, highlight Quit and press Enter. You should then reboot the system to ensure the new partition is made available.

However, it won't yet be accessible without a little more work. When Ubuntu is back up and running, open a GNOME Terminal window and issue the following command (again assuming that the new hard disk is `/dev/hdb`):

```
sudo mkfs.ext3 /dev/hdb1
```

Note that you need to specify the partition number in this instance. Because there's only one partition on the disk, this is number 1. Had you created two or more partitions, each would be numbered 1, 2, and 3 consecutively.

You've created an ext3-formatted partition, because this is the preferred standard used within Ubuntu. Other versions of Linux might use different file systems, such as ReiserFS.

Configuring Ubuntu to Use the Drive

Now the new drive is ready for use, but you need a way of making it available within the Ubuntu file system. Therefore, you need to create a mount point and also configure the system so that the disk is mounted automatically at boot.

As discussed in Chapter 14, creating a mount point is simply a matter of creating an empty folder. Therefore, you can create a directory in the root of the hard disk (or anywhere else) and call it something like `second_disk`. This directory must then be made writable, as follows:

```
sudo mkdir /second_disk
sudo chmod a+w /second_disk
```

Then you must edit the `/etc/fstab` file in order to make the new disk mount automatically. All you need to do is add a line at the end of the file, such as this:

```
/dev/hdb1 /second_disk ext3 default 0 2
```

Note that it's important that you add a carriage return (press Enter) after the line.

You can test your new hard disk by rebooting. When Ubuntu returns, you should find that the new disk is available by accessing the `/second_disk` directory. You can check its capacity by typing `df -h`.

Summary

In this chapter, we looked at streamlining your installation of Ubuntu. This involved speeding up the boot procedure by decreasing the boot menu delay and deactivating various unnecessary run-level scripts that get loaded at boot time. We also looked at optimizing your hard disk settings to allow for greater efficiency in loading and saving files.

Additionally, we investigated prelinking programs so that they load faster, recompiling the kernel so that it's optimized for your system, freeing disk space by various means, and adding a second hard disk.

In the next chapter, you'll learn how to perform backups to safeguard your data.



Backing Up Data

Every computer user knows that backing up data is vital. This is usually because every computer user has lost data at some point, perhaps because of a corrupted file or accidental deletion.

Some of the people behind Unix were highly aware of such occurrences, and built in several advanced and useful backup tools. These have been mirrored within Linux, with the result that creating and maintaining backups is easy.

In this chapter, we'll first look at what data should be backed up, and then explore two ways to make backups: using the SBackup utility and from the command line.

What Data Should You Back Up?

Data on your system can be classified into three broad types: program data, configuration data, and personal data. It's traditionally reasoned that backing up all types of data is inefficient and difficult, largely because it would mean backing up practically the entire hard disk. Because of this, you usually want to back up the latter two types of data: configuration and personal. The theory is that if your PC is hit by a hard-disk-wrecking disaster, you can easily reinstall the operating system from the CD/DVD. Restoring your system from backup is then simply a matter of ensuring the configuration files are back in place, so your applications work as you would like them to, and making sure that your personal data is once again made accessible.

Practically all the personal configuration data for programs you use everyday, as well as your personal data, is stored in your `/home` folder (although the configuration files for software used system-wide are stored in the `/etc` folder). If you take a look in your `/home` directory, you might think that previous sentence is incorrect. On a freshly installed system, the directory appears largely empty. However, most, if not all, of the configuration files are hidden; their directory and filenames are preceded with a period (`.`), which means that Linux doesn't display them during a standard directory listing.

To view hidden files and folders in the Nautilus file manager, select **View** ► **Show Hidden Files**. This can be quite an eye-opener when you see the masses of data you didn't even realize were there, as shown in the example in Figure 32-1. To view hidden files at the shell prompt, simply use the `-a` command option with the `ls` command:

```
ls -a
```

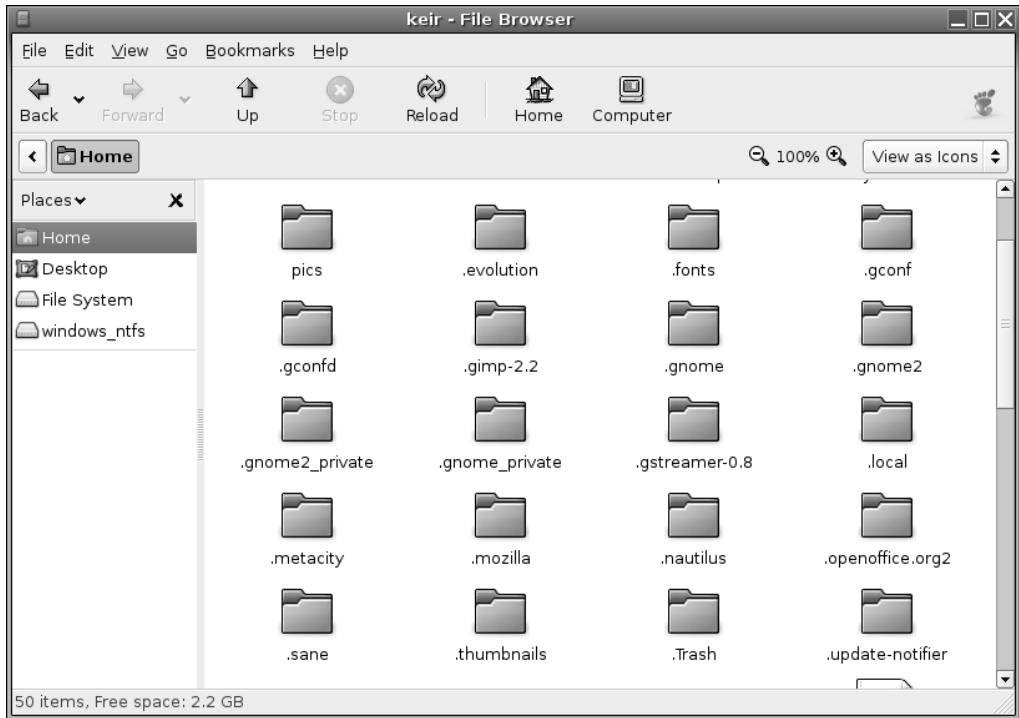



Figure 32-1. Most of the configuration files for programs are hidden—literally—in your /home folder.

The configuration files held in your /home folder relate solely to your user account. Any other users will have their own configuration files, entirely independent of yours. In this way, all users can have their own configuration settings for various applications, which can be backed up independently.

Under Ubuntu, you can back up both configuration data and personal files using Simple Backup Suite (SBackup), which can be downloaded from the Ubuntu software repositories.

Keep in mind that there's little point in making backups if you leave the resultant archive files on your hard disk. For full backup protection, the archives should be stored elsewhere, such as on an external hard disk, network mount or CD/DVD-ROM. Consider using GNOME CD/DVD Creator (click [Go](#) ► [CD/DVD Creator](#) on the menu of any open Nautilus window).

Using SBackup

Simple Backup Suite, also known as SBackup, is a series of programs that enables the quick and easy backup and also restoration of files. SBackup can back up personal data as well as system configuration files, and can store the resultant backup file on your hard disk (or a network mount attached to it), or in a remote Internet location, such as an FTP server.

SBackup was created courtesy of the Google Summer of Code sponsorship scheme and was designed with the help of Ubuntu developers. It's a new member of Ubuntu software

family, and it's likely that new features will be added all the time, so be sure to regularly update your system (see Chapter 9).

To install SBackup, open Synaptic Package Manager (System ► Administration), and then search for `sbackup`. Click its entry in the list of results, mark it for installation, and click Apply. Log out and back in again in order to update the system menus. You'll then find entries for the backup and restoration components of SBackup on the System ► Administration menu.

Backing Up Data via SBackup

To configure a backup, select System ► Administration ► Simple Backup Config. You'll see the Backup Properties dialog box, as shown in Figure 32-2. Using this dialog box, you can choose the files that SBackup backs up, as well as when it does so.

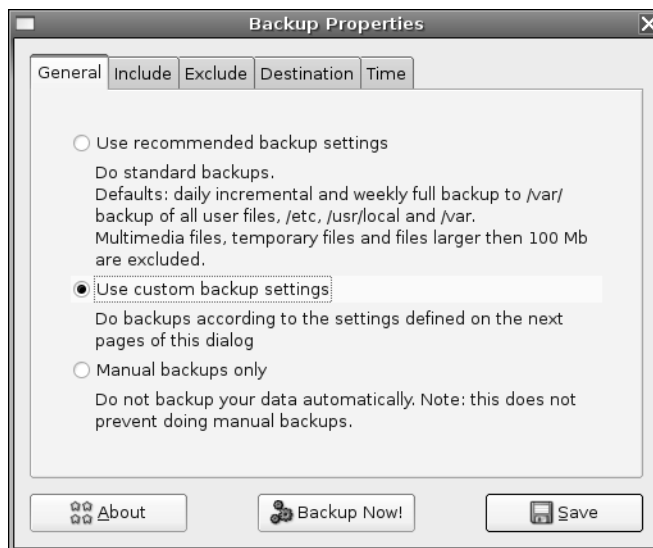


Figure 32-2. SBackup can work automatically or with custom settings you specify.

As listed on the General tab of the Backup Properties dialog box, SBackup can operate in three different modes:

Recommended backup: This is by far the best choice for fuss-free operation. SBackup will perform a daily backup of your `/home` folder, as well as the vital system data held in `/etc`, `/usr/local`, and `/var`. However, it will deliberately exclude any multimedia files (because of their large size), along with any temporary files and files of any type that exceed 100MB (again for size reasons). By default, the backup file created is placed in `/var/backup`.

Custom backup: This is effectively the same as the recommended backup, and includes the same list of file inclusions and exceptions, but you are able to edit the settings manually. For example, you might choose to include MP3 files, rather than excluding them, as with the recommended backup. The custom backup option lets you alter where the eventual backup file is saved and also lets you alter the time when the backup is made.

Manual backups: This effectively deactivates SBackup, so that it no longer periodically backs up files. However, you can still click the Backup Now! button to manually perform a backup according to the settings on the other Backup Properties dialog box tabs.

Note SBackup doesn't create a new backup file each time it runs because that would take too long. Instead, it creates an *incremental backup*, which means that it updates the last backup with files that have changed or been created since. Any files that haven't been updated since the last backup are unchanged.

Including Files and Folders in the Backup Job

Assuming that you've chosen a custom backup, and so are able to alter the backup settings, clicking the Include tab allows you to specifically define directories and files that you wish to include in the backup. Simply click the Add File or Add Directory button, and then browse to the relevant location (to add a directory, you'll need to click to open it before clicking the Open button).

Bear in mind that adding a directory does so recursively, which means that any directories contained within that directory are also backed up. For this reason, you don't need to specifically add your `/home/<username>` folder, because the entire `/home` folder is included in the backup by default. This means the backup will also include all other users' directories within `/home`, too.

Excluding Files and Folders From the Backup Job

You have a wide range of choices when it comes to excluding files and folders from the list. While directories can be excluded based on location, files can be excluded based on location, type of file, or size.

Clicking the Exclude tab reveals a set of side tabs on the left side of the program window, which allow you to exclude items from the backup as follows:

Paths: To exclude a specific file or folder, click this side tab. As with including files, click the Add File or Add Directory button, and then browse to the relevant location.

File Types: To exclude certain types of files, click this side tab, as shown in Figure 32-2. After clicking the Add button, you'll see that you can choose from a brief list of standard file types or filter by file extension (such as `.mp3` for MP3 files or `.zip` for compressed Zip files).

Regex: If you're competent at using regular expressions, as outlined in Chapter 15, you can use them to specify extremely precise rules by clicking this side tab.

Max Size: Any files larger than the stated size on this tab aren't backed up. By removing the check next to the Do Not Backup line, you can deactivate this feature (although that could lead to massive backup files).

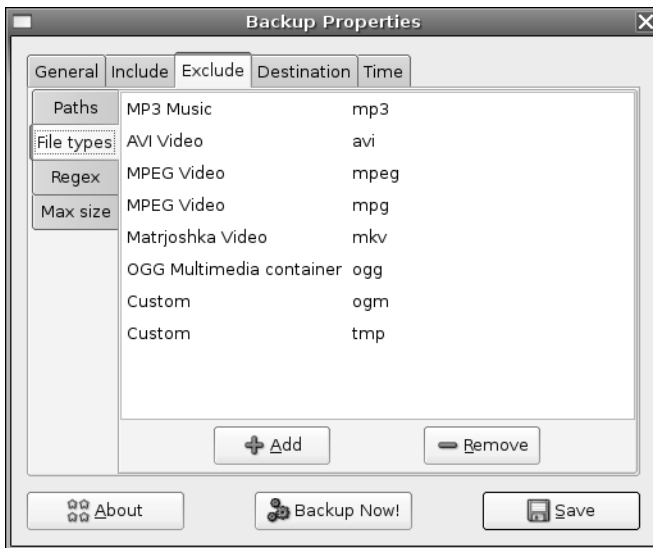


Figure 32-3. Excluding certain types of large files will lead to far smaller backup files.

Changing the Backup File Destination

By default, the backup file created by SBackup is stored in the `/var/backup` directory. By clicking the Destination tab, you can choose to save it in a different location on your hard disk. Of course, if you have a network mount, you can also opt to save it there.

Note Remember that there's little point in retaining backups on your hard disk. They should be copied to a removable storage device, such as a CD-R/RW disc.

You can even transfer the backup file across the Internet via SSH file transfer or the older FTP standard. To do so, simply enter the protocol, username, password, and URL in the following format:

```
ftp://username:password@myftpsite.com
```

It's important to precede the address with the protocol you intend to use, whether it's SSH or FTP.

Changing the Time Period Between Backups

Clicking the Time tab lets you set the frequency of the backup. You can opt to back up hourly, daily, weekly, or monthly. You can also set the exact time of the backup if necessary. For example, you could set a backup to take place every week on a Tuesday at 12.30 p.m. Simply select the interval period from the Do Backups drop-down list, and then select from the Day of Month, Day of Week, Hour, and Minute lists, as necessary.

You can also elect to perform a full backup after a certain number of days have passed (up to 1,000). A full backup means that SBackup creates a new backup file, rather than appending an already existing file with any new or edited files.

Restoring Data via SBackup

If the worst happens and you need to restore any number of files from the backup, you can click System ► Administration ► Simple Backup Restore.

The first step is to click in the Available Backups drop-down list to choose a backup file from which to restore. The filenames contain the dates and times the backups were made, and it makes sense to choose the latest file (unless you want to revert to an older version of a file).

Once the backup has been selected, the files it contains will be displayed beneath. Each directory will have a small triangle to its left, which you can click to expand the directory to show its contents.

After you've found the file or directories you want to restore, highlight them and then click the Restore button. But beware, because this will rewrite the files and directories to their original locations—files or directories already there with matching filenames will be overwritten!

If you want to restore the files to a different location, click the Restore As button, and then choose a folder.

Caution Restored files and directories are owned by root. This is because SBackup runs with superuser powers. Therefore, one of the first things you'll have to do is use the `chown` command, preceded by `sudo`, to change the ownership and group of the file to what they were originally. See Chapter 14 for more details about file ownership and how the `chown` command works.

Making Backups from the Command Line

Although SBackup allow the uninitiated to make quick backups, the `tar` program is preferred by Linux old-timers. This creates `.tar` files and is one of the original carryovers from Unix. `tar` stands for Tape ARchive and refers to backing up data to a magnetic tape backup device. Although `tar` files are designed for backup, they've also become a standard method of transferring files across the Internet, particularly with regard to source files or other installation programs.

A `tar` file is simply a collection of files bundled into one. By default, the `tar` file isn't compressed, although additional software can be used to compress it. `tar` files aren't very sophisticated compared to modern archive file formats. They're not encrypted, for example, but this can also be one of their advantages.

Note Linux comes with a couple more backup commands, which you might choose to use. They are `cpio` and `pax`. Both aim to improve on `tar` in various ways, but neither is broadly supported at the moment. `cpio` is installed by default under Ubuntu, and `pax` can be found via the Synaptic Package Manager. Examine their man pages for more details.

Creating tar Files

Perhaps unsurprisingly, tar files are created at the console using the tar command. Usually, all that's needed is to specify a source directory and a filename, which can be done like so:

```
tar -cf mybackup.tar /home/keir/
```

This will create a backup called `mybackup.tar` based on the contents of `/home/keir/`. `tar` is automatically recursive so, in this example, it will delve into all subdirectories beneath `/home/keir/`. The `-c` command option tells tar you're going to create an archive, and the `-f` option indicates that the filename for the archive will immediately follow. If you don't use the `-f` option, tar will send its output to standard output, which means that it will display the contents of the archive on the screen.

If you typed in a command like the preceding example, you would see this message:

```
Removing leading '/' from member names.
```

This means that the folders and files added to the archive will all have the initial forward slash removed from their paths. So, rather than store a file in the archive as this:

```
/home/keir/Mail/file1
```

The file will be stored as follows:

```
home/keir/Mail/file1
```

The difference between the two forms concerns when the files are later extracted from the archive. If the files have the initial slash, tar will write the file to `/home/keir/Mail/file1`. If there's already a file of that name in that location, it will be overwritten. On the other hand, with the leading slash removed, tar will create a new directory wherever you choose to restore the archive. In this example, it will create a new directory called `home`, and then a directory called `keir` within that, and so on.

Because of the potential of accidentally overwriting data by specifying absolute paths in this way, a better way of backing up a directory is simply to change into its parent and specify it without a full path:

```
cd /home/  
tar -cf mybackup.tar keir
```

When this particular archive is restored, it will simply create a new folder called `keir` wherever it's restored.

Compressing tar Archives

You can also compress the archive from within tar, although it actually calls in outside help from either `bzip2` or `gzip`, depending on which you specify.

To create a tar archive compressed using `bzip2`, the following should do the trick:

```
tar -cjf mybackup.tar.bz2 keir
```

This will create a compressed backup from the directory `keir`. The `-j` command option passes the output from `tar` to the `bzip2` program, although this is done in the background. Notice the change in the backup filename extension to indicate that this is a `bzip2` compressed archive.

The following command will create an archive compressed with the older `gzip` compression:

```
tar -czf mybackup.tar.gz keir
```

This uses the `-z` command option to pass the output to `gzip`. This time, the filename shows it's a `gzip` compressed archive, so you can correctly identify it in the future.

Extracting Files from a tar Archive

Extracting files using `tar` is as easy as creating them:

```
tar -xf mybackup.tar
```

The `-x` option tells `tar` to extract the files from the `mybackup.tar` archive.

Extracting compressed archives is simply a matter of adding the `-j` or `-z` option to the `-x` option:

```
tar -xjf mybackup.tar.bz2
```

Note Technically speaking, `tar` doesn't require the preceding hyphen before its command options. However, it's a good idea to use it anyway, so you won't forget to use it with other commands in the future.

Viewing tar Archive Information

To view the contents of a `tar` archive without actually restoring the files, use the `-t` option:

```
tar -tf mybackup.tar |less
```

This example adds a pipe into `less` at the end, because the listing of files probably will be large and scroll off the screen. Just add the `-j` or `-z` option if the `tar` archive is also compressed.

In addition, you can add the `-v` option to all stages of making, extracting, and viewing an archive to see more information (chiefly the files that are being archived or extracted).

Typing `-vv` provides even more information:

```
tar -cvvf mybackup.tar keir
```

This will create an archive and also show a complete directory listing as the files and folders are added, including permissions.

Saving the File to a CD-R/RW

Once the tar file has been created, the problem of where to store it arises. As I mentioned earlier, storing backup data on the same hard disk as the data it was created to back up is foolish, since any problem that might affect the hard disk might also affect the archive. You could end up losing both sets of data!

If the archive is less than 700MB, it might be possible to store it on a CD-R or CD-RW. To do this from the command line, first the file must be turned into an ISO image, and then it must be burned.

To turn it into an ISO image, use the `mkisofs` command:

```
mkisofs -o backup.iso mybackup.tar.bz2
```

You can then burn the ISO image to a CD by using the `cdrecord` command. Before using this, you must determine which SCSI device number your CD-R/RW drive uses (all CD-R/RW or DVD-R/RW drives are seen as SCSI devices, even if they're not). Issue the following command (with superuser powers—`cdrecord` accesses the system hardware directly):

```
sudo cdrecord -scanbus
```

You should find the device numbers listed as three numbers separated by commas. To burn the backup image, all you need to do is enter a command in this format:

```
sudo cdrecord dev=<dev number> speed=<speed of your drive> mybackup.iso
```

On a typical system, this might take the following form:

```
sudo cdrecord dev=0,0,0 speed=24 mybackup.iso
```

Note Remember that you can use the `man` command to learn about any commands, including those detailed here. In the case of `cdrecord` in particular, this is worthwhile, because it's a very powerful piece of software.

Summary

In this chapter, we looked at making backups. First, you saw how to verify where your personal and other vital data is stored. Then we looked at how the `SBackup` tool can be used to back up system configuration and personal data. Finally, you learned how to use `tar` at the command line to back up any kind of data.

In the next chapter, we'll look at how tasks can be scheduled to occur at various times under Ubuntu.



Scheduling Tasks

In this book, you've learned about various tasks you can perform to keep Ubuntu running smoothly. You may decide that you want some of these tasks to occur on a regular basis. For example, perhaps you want your `/home` folder to be backed up every day, or perhaps you want to clean the `/tmp` folder to ensure that you always have enough free disk space. You could carry out each task individually, but human nature would no doubt step in, and you would forget, or you might perform the action twice, because you've forgotten that you've already done it.

As you might expect, Linux is able to automate the running of particular tasks. They can either be run periodically at scheduled times or as one-time tasks. Using Linux's scheduling features is explained in this chapter.

Scheduling with crontab

Under Ubuntu, the main way of scheduling tasks is via the `cron` daemon. This works on behalf of the user in order to schedule individual tasks, and it is also used by the system to run vital system tasks, although a different way of working is used in each case.

For `cron` to run user-scheduled tasks, it reads a file called `crontab`. Each user has her own version of this file, which is stored in the `/var/spool/cron/crontabs` directory. This file can be edited in a text editor, but a special command should be used to do so.

Note System-wide tasks are handled by the `/etc/crontab/` file. This runs scripts contained in `/etc/cron.hourly`, `/etc/cron.daily`, and so on, depending on when the tasks are meant to be run (every hour, day, week, or month). The average user never needs to bother with system-wide `cron` jobs. These are handled by the internal system, and programs create their own entries as and when necessary.

The `cron` daemon starts at bootup and simply sits in the background while you work, checking every minute to see if a task is due. As soon as one comes up, it commences the task, and then returns to a waiting status.

Creating a Scheduled Task

Adding a scheduled task is relatively easily and is done via the shell. Entering the following command will cause your personal crontab file to be loaded into the GNU nano text editor, ready for editing:

```
crontab -e
```

If this is the first time you've edited your crontab file, it will most likely be completely empty (in fact, you might see a message about crontab having to create the file from scratch). However, don't be put off. Adding a new entry is relatively easy and normally takes the form of something like this:

```
01 12 15 * * tar -cjf /home/keir/mybackup.tar.bz2 /home/keir
```

Let's examine the line piece by piece. The first part—the numbers and asterisks—refers to when the task should be run. From left to right, the fields refer to the following:

- Minutes, from 0 to 59
- Hours, in 24-hour time, so from 0 to 23
- Day dates, for the day of the month, from 1 to 31 (assuming the month has that many days)
- Months, from 1 to 12
- Day, for a particular day, either from 0 to 6 (0 is Sunday), or specified as a three-letter abbreviation (mon, tue, wed, and so on)

In the example, the task is set to run at the first minute at the twelfth hour (midday) on the fifteenth day of the month. But what do the asterisks stand for? They're effectively wildcards and tell cron that every possible value applies. Because an asterisk appears in the month field, this task will be run every month. Because an asterisk appears in the day field, the task will be run every day.

You might have noticed a logical contradiction here. How can you specify a day if you also specify a date in the month? Wouldn't this seriously limit the chances of the task ever running? Yes, it would. If you were to specify sat, for example, and put 15 in the date field, the task would run on only the fifteenth of the month if that happened to be a Saturday. This is why the two fields are rarely used in the same crontab entry, and an asterisk appears in one if the other is being used.

After the time and date fields comes the command itself: tar. As you learned in the previous chapter, tar is designed to back up your personal data.

Only standard BASH shell commands can be used in the command section. cron isn't clever enough to interpret symbols such as the tilde (~) as a way of referring to your home directory. For this reason, it's best to be very thorough when defining a cron job and always use absolute paths.

Let's take a look at another example (shown in Figure 33-1):

```
59 23 * * 0-3 tar -cjf /home/keir/mybackup.tar.bz2 /home/keir
```

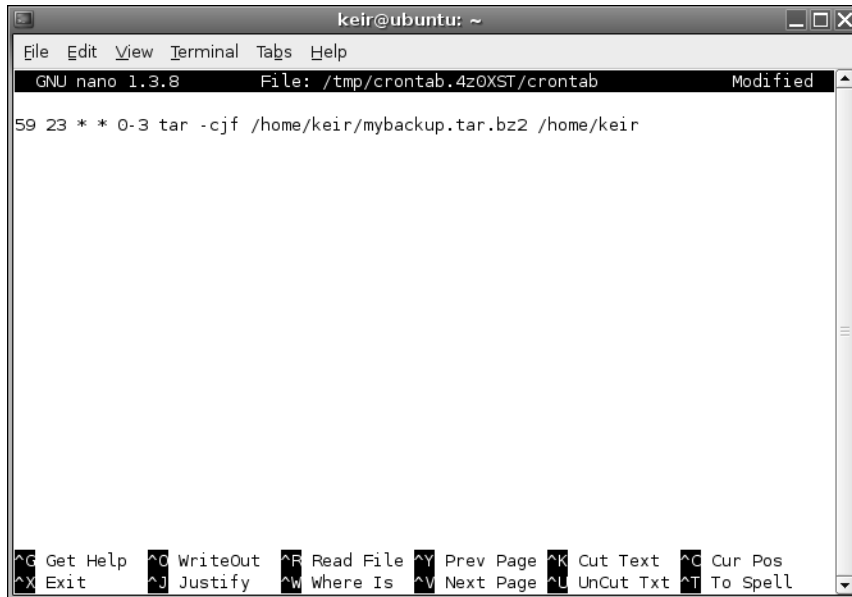


Figure 33-1. *Editing crontab lets you schedule tasks using the nano text editor.*

The first field says that this task will run at the fifty-ninth minute of the twenty-third hour (that is, one minute before midnight). The date and month field have asterisks, so this implies that the task should run every day and every month. However, the day field contains 0-3. This says that the task should run on only days 0 through to 3, or Sunday through Wednesday.

You can have as many cron entries as you like; simply give each a separate line. You don't need to put them in date or time order. You can just add them as and when you see fit.

When you're finished, save the file and quit GNU nano in the usual way (by pressing Ctrl+X).

Editing crontab with the GUI

Although it's not installed by default, the `gcrontab` program lets you edit the crontab file via a graphical interface. You'll find the program in the Ubuntu software repositories and can install it using the Synaptic Package Manager (see Chapter 8).

Using the program is similar to manually editing crontab. Start by clicking the Add New button to add a scheduled job, and then use the various icons on the second toolbar to alter the time, day, month, and so on. You can then click the lightning icon to add the actual command.

Note To be frank, I can't see any reason to use `gcrontab` over simply editing the crontab file manually, but if you really hate using the command line, it might be worth investigating.

Scheduling with anacron

If cron has an Achilles' heel, it is that it expects your computer to be up and running all the time. If you schedule a task for around midnight, as in one of the previous examples, and your computer isn't switched on at that time, the task simply won't run.

anacron was created to fix this problem (see Figure 33-2). It also can run scheduled tasks, but unlike cron, it doesn't rely on exact times or dates. Instead, it works on the principle of time periods. For example, tasks can be set to run every day. In fact, tasks can be set to run every x number of days, regardless of whether that's every two days or every hundred thousand. It also doesn't matter if the computer is shut down and rebooted during that time; the task will be run only once in the specified time period. In addition, tasks can be specifically set to run at the beginning of each month, regardless of the length in days of each month.

```

keir@ubuntu: ~
File Edit View Terminal Tabs Help
# /etc/anacrontab: configuration file for anacron
# See anacron(8) and anacrontab(5) for details.

SHELL=/bin/sh
PATH=/sbin:/bin:/usr/sbin:/usr/bin

# These replace cron's entries
1 5 cron.daily nice run-parts --report /etc/cron.daily
7 10 cron.weekly nice run-parts --report /etc/cron.weekly
@monthly 15 cron.monthly nice run-parts --report /etc/cron.monthly
~
~
~
~
~
~
9,1 All

```

Figure 33-2. *anacron* is used to run tasks periodically, such as every couple of days.

anacron is primarily designed to be used for system maintenance, and the `anacrontab` file holds the details of the tasks. Unlike as with `crontab`, each user doesn't have his own `anacrontab` file. However, there's no reason why you can't add your own commands to the main `anacrontab` file. This file can be edited in any text editor, and you don't need to use a special command to edit it (as with `crontab`), although you'll need to adopt superuser powers to do so.

Each line in `anacrontab` takes the following form:

```
days    delay    name of task    command
```

The `days` field holds the number of days in between the running of the task. To set the task to run every day, you would enter 1. To make the task run every nine days, you would add 9. To set it to run monthly, you should type `@monthly`.

The delay field tells `anacron` how long to wait after it's run before running the task, specified in minutes. It is necessary because `anacron` is usually run at bootup, and if it were to run the scheduled tasks instantly, the machine would grind to a halt because it is already busy. A delay of five minutes is usually adequate, although if some tasks are already scheduled to run on the same day before that task, you should allow enough time for them to finish.

The name of task field is for your personal reference and shouldn't contain either slashes or spaces (hint: separate words using underscores or full stops).

The command field is, as with `crontab`, the shell command that should be run.

Note `anacron` is run as the root user, so if you do add your own entry to `anacrontab`, any files it creates will be owned by root, too. If you use `anacron` to create a backup of your `/home` directory, for example, the resultant backup file will be owned by root, and you'll need to use the `chown` command to change its ownership so you can access it. See Chapter 14 for more information about the `chown` command.

Let's look at an example of an `anacrontab` entry:

```
1 15 backup_job tar -cjf /home/keir/mybackup.tar.bz2 /home/keir
```

This will run the specified `tar` command every day (because 1 is in the days field), and with a delay of 15 minutes after `anacron` is first run.

`anacron` is run automatically every time you boot, but you can also run it manually by simply typing it at the command prompt (with superuser powers):

```
sudo anacron
```

Using `at` to Schedule One-Off Tasks

What if you quickly want to schedule a one-time-only task? For this, you can use the `at` command.

Adding a job with `at` is very easy, largely because the `at` command accepts a wide variety of time formats. For example, typing the following at the command prompt will run a job at lunchtime tomorrow:

```
at noon tomorrow
```

It really is as simple as that!

Alternatively, you can specify a time, date, and even a year:

```
at 13:00 jun 25 2008
```

This will run the job at 1 p.m. on June 25, 2008. The various time and date formats are explained in the `at` command's man page.

Once the `at` command containing the date has been entered, you'll be presented with a mock shell prompt. Here, you can type the commands you want to run. Many shell commands can be entered, one after the other; just press Enter between them. Then press Ctrl+D to signal that you're finished editing. At this point, `at` will confirm the time and write the task into its list.

You can view the list at any time by typing `atq`. This will show a list of numbered jobs. You can remove any job by typing `atrm`, followed by its `atq` job number. For example, the following will remove the job numbered 9 in the `atq` list:

```
atrm 9
```

Summary

In this brief chapter, we looked at how you can schedule tasks under Ubuntu, which is essentially making programs run at certain times. We examined the `crontab` and `anacron` facilities, which can schedule tasks to run periodically, and we also examined the `at` service, which can schedule one-off tasks to run at certain times.

In the final chapter of this book, we will look at how you can access your Ubuntu computer remotely—theoretically, from any Internet-equipped location in the world.



Accessing Computers Remotely

One area where Linux particularly excels is in its support for networking, including the Internet. If you wish to learn about how networks operate on a fundamental level, then Linux is an ideal choice, because it puts you in virtually direct contact with the technology.

The widespread integration and support for networking extends to several useful system tools, which let you access Linux across any kind of network, including the Internet. In fact, it's even possible to access a Linux machine running on a different continent, just as if you were sitting in front of it!

This chapter looks at the many ways you can access your Ubuntu computer remotely. In addition, we look at the ways that you can use Ubuntu to access almost any other computer, including Windows XP Professional PCs.

Using Secure Shell

The history of Unix has always featured computers connecting to other computers in some fashion, whether they were dumb terminals connecting to a mainframe computer or Unix machines acting as nodes on the fledgling Internet. Because of this, a wide variety of techniques and protocols were invented to allow computers to communicate and log in to each other across networks. However, while these still work fine over the modern Internet, we're now faced with new threats to the privacy of data. In theory, any data transmitted across the Internet can be picked up by individuals at certain key stages along the route. If the data isn't protected in any way, it can be easily intercepted and read.

To counter such an occurrence, the `ssh` suite of programs was created. Although these programs started as open source, they gradually became proprietary. Therefore, several newer open-source versions were created, including the one used on the majority of Linux distributions (including Ubuntu): OpenSSH.

The goal of `ssh` is to create a secure connection between two computers. You can then do just about any task, including initiating a shell session so you can use the remote computer as if you were sitting in front of it, or copying files to and from `ssh` uses various techniques at both ends of the connection to encrypt not only the data passing between the two machines, but also the username and password.

Note This chapter refers to *remote* and *local* machines. The *remote* machine is the computer you're connecting to across the network or Internet. The *local* machine is the one you're actually sitting in front of. These two terms are widely used within documentation describing networking.

Logging In to a Remote Computer

The most basic type of `ssh` connection is a remote login. This will give you a command prompt on the remote computer, as if you had just sat down in front of it and opened GNOME Terminal.

But before you can log in to any machine via `ssh`, you'll need to be sure the remote computer is able to accept `ssh` connections. This means that it needs to be running the `ssh` server program (referred to as a *service*), and also that its firewall has an open port for incoming connections.

The two major components of OpenSSH are the client and server. Most distributions install both items and run the server component all the time. However, only the client component of SSH is installed under Ubuntu. To install the server component, and therefore access your Ubuntu system remotely, you'll need to open Synaptic Package Manager (System Administration) and search for `openssh-server`. Click to install it. Configuration will be automatic, although if you're using the Ubuntu firewall (see Chapter 9), you will need to configure an incoming rule to open port 22.

Tip If you use Firestarter, as described in Chapter 9, you can simply select the default incoming SSH rule. There's no need to manually specify a port number.

Initiating an `ssh` remote shell session with a remote machine is usually achieved by typing something similar to the following at a command prompt on the local machine:

```
ssh <username>@<IP address>
```

In other words, you specify the username you want to log in as, as well as the IP address of the machine. If there's a fully qualified domain name (FQDN) for the system you want to access, you could specify that instead of the IP address.

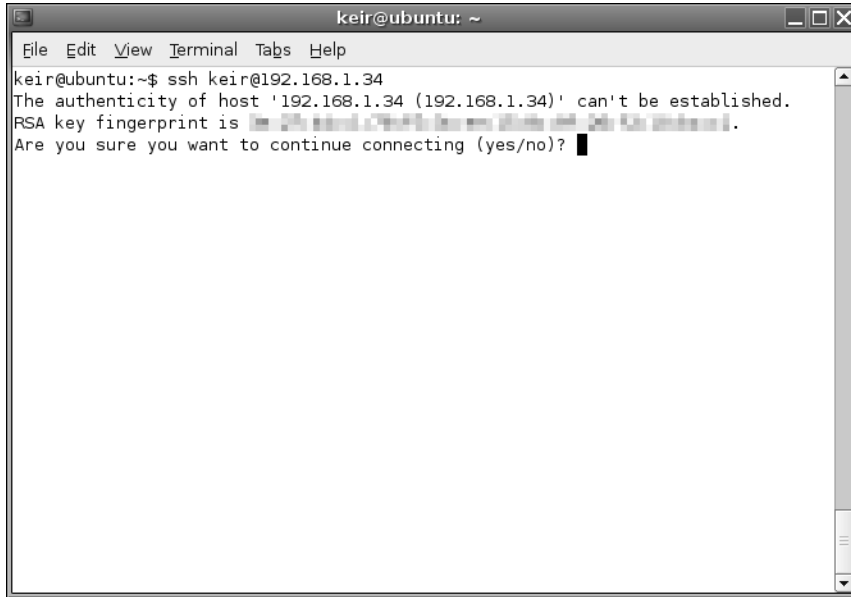
Note An FQDN is the hostname of a system plus its Internet address, such as `mycomputer.example.com`. Unless you have had this function specifically set up for you by a system administrator, you'll probably have to connect via IP addresses. However, if you rent a web server, you might be able to `ssh` into it using the domain name of the server.

You'll be prompted for your password which, obviously, is the password for the account you're trying to log in to on the *remote* computer.

When you log in for the first time, you'll see the following message:

The authenticity of the host <host IP address> can't be established

Figure 34-1 shows an example. This means that the remote computer's encryption key hasn't yet been added to your PC's store file. However, once you agree to the initial login, the encryption key will be added, and it will be used in the future to confirm that the remote computer you're connecting to is authentic.



```
keir@ubuntu: ~  
File Edit View Terminal Tabs Help  
keir@ubuntu:~$ ssh keir@192.168.1.34  
The authenticity of host '192.168.1.34 (192.168.1.34)' can't be established.  
RSA key fingerprint is [blurred].  
Are you sure you want to continue connecting (yes/no)? █
```

Figure 34-1. When logging in via ssh for the first time, you'll need to confirm acceptance of the encryption key.

Note There's a fine line between security concern and paranoia. Connecting to a remote machine for the first time and accepting its ssh key is considered insecure by some people, because you cannot be 100% sure that the remote machine is the one you want to connect to. It might have been swapped for a different machine by hackers (or some such theory). In addition, the key might be intercepted on its journey to you. Because of this, those who are highly security conscious can use the ssh-keygen program to create a key on the remote machine first, and then import it to the local machine via floppy disk before logging in. See the ssh-keygen man page for more details.

After confirming that you want to make the connection, you'll be invited to enter the password for the user account under which you initiated the ssh connection. Once this is done, you should find yourself with a shell login on the remote computer. You can run the same commands as usual and perform identical tasks.

The machine you're logged in to will show no symptoms of being used remotely. This isn't like the movies, where what you type on your local machine is somehow mirrored on the remote machine for all to see. However, obviously, if a user of the remote machine were to view

her network connections using something similar to the `netstat` command, then she would see another computer attached via `ssh`.

To end an `ssh` session, simply type `exit`. This will then return you to the command prompt on your own machine.

Tip There's a version of the `ssh` client that runs on a variety of non-Linux operating systems, making it possible to log in to your Ubuntu machine from a Windows computer. The program is called PuTTY and can be downloaded from www.chiark.greenend.org.uk/~sgtatham/putty/.

MANAGING REMOTE SESSIONS

Whenever you open any kind of shell to enter commands and run programs, you might have noticed that any commands you start running last only as long as the shell window is open. When the shell window is closed, any task running within it ends, too. This is because the shell is seen as the “owner” of the process, and when the owner dies, any processes it started also die.

When using `ssh` to start a remote shell session, this also applies. Whenever you log out, any tasks you were running are ended. This can be annoying if, for example, you've started a lengthy download on the remote machine. Effectively, you must remain logged in via `ssh` until the download has finished.

To get around this, you can use the handy `screen` program. This isn't specifically designed to be an aid to remote logins, but there's no reason why it cannot be used in such a situation.

The `screen` program effectively starts shell sessions that stick around, even if the shell window is closed or the `ssh` connection is ended or lost. After logging in to the remote computer via `ssh`, you can start a `screen` session by simply typing the program name at the prompt:

```
screen
```

After pressing the spacebar as prompted to start the program, there won't be any indication that you're running a `screen` session. There's no taskbar at the bottom of the terminal window, for example. `screen` works completely in the background.

Let's consider what happens when you detach and then reattach to a `screen` session. To detach from the `screen` session, press `Ctrl+A` and then `Ctrl+D`. You'll then be returned to the standard shell and, in fact, you could now disconnect from your `ssh` session as usual. However, the `screen` session will still be running in the background on the remote computer. To prove this, you could log back in, and then type this:

```
screen -r
```

This will resume your `screen` session, and you should be able to pick up quite literally where you left off; any output from previous commands will be displayed.

To quit a `screen` session, you can either type `exit` from within it or press `Ctrl+A`, and then `Ctrl+\` (backslash).

The `screen` program is very powerful. To learn more about it, read its man page. To see a list of its keyboard commands, press `Ctrl+A`, and then type a question mark (?) while `screen` is running.

Transferring Files Between Remote Computers

The `ssh` utility brings with it two basic ways of transferring files between machines: `scp` and `sftp`. `scp` is fine for smaller file transfers, but if you want to copy a lot of files, `sftp` is probably a better choice.

Using `scp`

Strictly speaking, `scp` is merely a program that copies files from one computer to another in a secure fashion using the underlying `ssh` protocol. You don't have to be logged in to another computer via `ssh` to use it. For example, if I were merely browsing my own computer and wanted to transfer a file to a remote computer, I could type:

```
scp myfile keir@<IP address>:/home/keir/
```

I would replace `<IP address>` with the IP address of the computer to which I wanted to send the file. In other words, you must first specify the local file you want to copy across, and then provide the login details for the remote computer in the same format as with an `ssh` login. Then, after a colon, you specify the path on the *remote* computer where you would like the file to be copied.

Note If it helps, consider the latter part of the `scp` command after the filename as one large address: first you provide your username, then the computer address, and then the path.

Using the command when you *are* logged in to another computer via `ssh` works in exactly the same way. Let's consider an example.

Assume there are two computers: A and B. I have a user account on each one. So sitting at the keyboard of A, I establish an `ssh` connection with B by typing the following:

```
ssh keir@computer_B
```

This lets me log in to B as if I were sitting in front of it. I spot a file called `spreadsheet.xls` that I want to copy to my local machine (A). I therefore issue the following command:

```
scp spreadsheet.xls keir@computer_A:/home/keir/
```

This will copy the file from computer B to computer A and place it in the `/home/keir/` directory.

Tip With `scp`, you can copy entire directories, too. Simply add the `-r` command option, like so: `scp -r mydirectory <username>@<IP address>:/path/`.

Using sftp

To copy a lot of files to or from a remote computer, the `sftp` program is the best solution. If you've ever used a shell-based `ftp` program, you'll feel right at home, because `sftp` isn't very different.

The difference compared to `scp` is that you use `sftp` to establish a connection with the remote computer in a similar way to `ssh`, except you can use only certain commands. You can initiate a `sftp` session by using this command format:

```
sftp <username>@<IP address>
```

The same rules as when you're logging in with `ssh` apply, both in terms of formatting the login command and also confirming the encryption key if this is the first time you've logged in.

The `sftp` commands are fairly basic. For example, to copy a file from the remote machine, simply type this:

```
get <filename>
```

This will copy the file into the directory you were in on the local machine before you started the `sftp` session.

By specifying a path after the filename, the file will be copied to the specified local directory:

```
get spreadsheet.xls /home/keir/downloaded_files/
```

Sending files from the local machine to the remote machine is just as easy:

```
put <filename>
```

By specifying a path after the filename, you can ensure the file is saved to a particular remote path.

One useful thing to remember is that any command preceded by an exclamation mark (!, called a bang in Linux-speak) is executed on the local machine as a shell command. So, if you wanted to remove a file on the local machine, you could type this:

```
!rm -rf <filename>
```

Simply typing a bang symbol on its own starts a shell session on the local machine, so you can perform even more tasks. When you're finished, type `exit` to return to the `sftp` program.

For a list of popular `sftp` commands, see Table 34-1.

Table 34-1. *sftp* Commands

Command	Function
<code>cd</code>	Change the remote directory
<code>lcd</code>	Change the local directory
<code>get</code>	Download the specified file
<code>ls</code>	List the remote directory
<code>lls</code>	List the local directory
<code>mkdir</code>	Create a directory on the remote machine
<code>lmkdir</code>	Create a directory on the local machine

Table 34-1. *sftp* Commands

Command	Function
put	Upload the specified file to the remote machine
pwd	Print the current remote directory
rmdir	Delete the remote directory
rm	Delete the remote file
exit	Quit <i>sftp</i>
!command	Execute the specified command on the local machine
!	Start a temporary local shell session (type <i>exit</i> to return to <i>sftp</i>)
help	Show a list of commands

Accessing GUI Applications Remotely

So far, we've looked at connecting to a remote machine using command-line tools. But Ubuntu is based around the graphical desktop, so is there any way of running, say, a Nautilus file browser window so you can manipulate files on the remote machine? Yes!

The graphical subsystem of Linux, X, is designed to work across a network. In fact, if you run Linux on your desktop PC, X *still* works via a loopback network within your machine (meaning that network commands are sent out but addressed to the very same machine on which they originated). Because of this, it's possible to make programs on a remote machine run on a local machine's X server. The actual work of running the application is handled by the *remote* machine, but the work of displaying the graphics is handled by the *local* machine.

Caution X connections across a network can be a little slow and certainly not as snappy as running the same application on the local machine. This lag can become irritating after a while.

Running X Applications on a Remote Computer

Unfortunately, X server communications aren't normally encrypted, so if one machine were to simply connect to an X server over a network (or even the Internet), the data transfer would be unencrypted and open to eavesdroppers.

But *ssh* once again comes to the rescue. You can configure *ssh* so that X applications on the remote computer can be run on the local machine, with the data sent through the *ssh* connection. Log in to the remote machine using *ssh*, but also specify the *-X* flag:

```
ssh -X <username>@<IP address>
```

When you're logged in, you can simply start any application by typing its name as usual. The only difference is that the program will appear on the screen of the local machine, rather than on the remote machine, as shown in Figure 34-2.

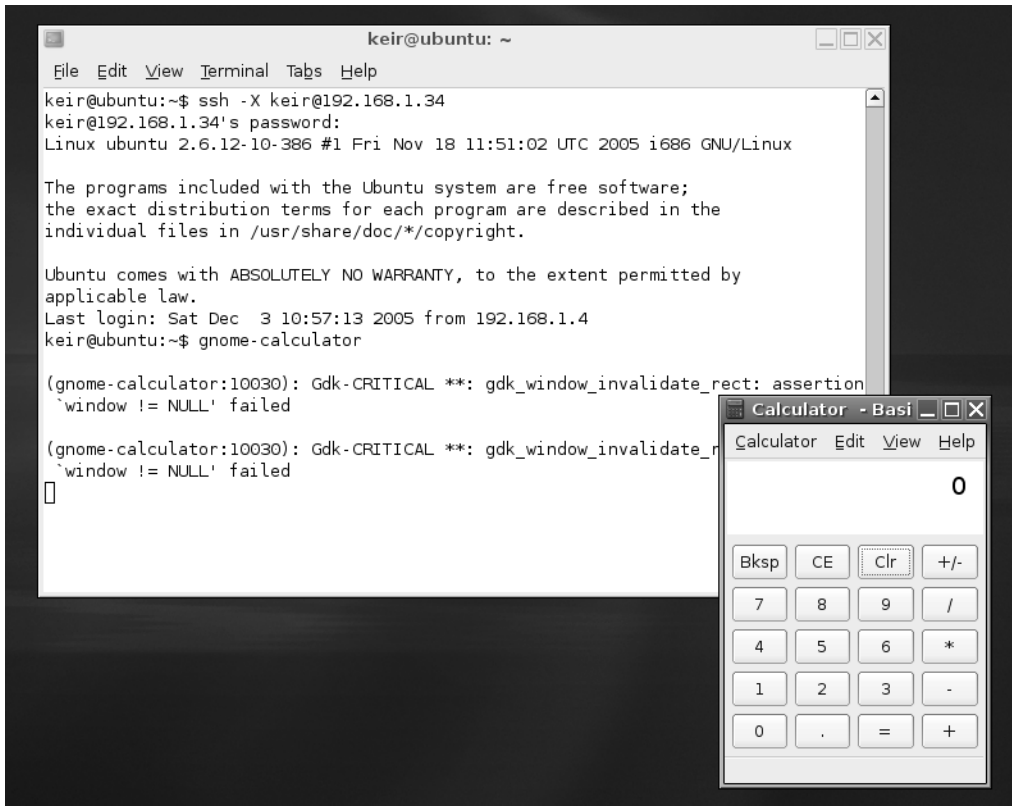


Figure 34-2. Although the calculator application appears on the local computer's display, it's actually running on the remote machine.

Using X across the Internet or even a local network isn't very fast, and you can expect delays when you open menus or if the screen must frequently redraw. However, it can prove very useful.

Running the GNOME Desktop Across a Remote Connection

It's even possible to run the entire GNOME desktop across an ssh connection. In other words, you can access the desktop of a remote computer as if you were sitting in front of it.

You can do this by starting a stripped-down X server shell on the local computer, and then initiating the GNOME desktop on the remote computer via an ssh connection.

Note If the remote computer is already running a desktop, yours will run in addition to it. Anyone using the remote computer won't be aware of what you're doing, unless they specifically look for evidence. However, sound isn't part of what X takes care of, so any sounds normally made during GNOME startup, such as the jingle, will be heard on the *remote* machine! The same applies for any error noises, such as beeps.

On the local machine from which you would like to initiate the connection, close all open programs and then kill the X server by pressing Ctrl+Backspace. If you find the GUI restarts, press Ctrl+Backspace again.

Log in at the command prompt and then type:

```
xinit
```

This will start a rudimentary X session, complete with an xterm shell window. You can then use the ssh command to log in to your remote machine, specifying that X communications should be allowed over the connection:

```
ssh -X <username>@<IP address>
```

And then type the following:

```
gnome-session
```

After a few seconds, the remote computer's GNOME desktop will appear. Make sure you don't close the original xterm window, because this owns the gnome-session process, so closing it would kill the desktop.

Once again, the remote desktop will be fairly slow to respond to mouse clicks and keyboard strokes, but it should be usable.

If you intend to do this often, you may wish to configure a command-line-only run level on the local machine. This will mean you boot to a non-GUI login prompt and can then initiate your own X session manually (either running GNOME on your local machine or on a remote machine). Most Linux distributions, such as Red Hat or SUSE, define run level 3 as command-line-only, so you can do the same.

Note Under Ubuntu, run levels 2 through 5 are exactly the same. Run level 2 is used all the time, and run levels 3 through 5 are effectively spares that can be altered however you wish.

Type the following to stop the GUI starting automatically under run level 3:

```
sudo rm -f /etc/rc3.d/S13gdm
```

Following this, you'll need to modify the /etc/inittab file so that you boot to the modified run level 3. The following will load it into the Gedit text editor:

```
sudo gedit /etc/inittab
```

Look for the line that begins `id:2:initdefault:` and change it so it reads `id:3:initdefault:`. Then save the file and reboot.

Once you've logged in, you'll have the choice of either starting the X server shell for a remote connection, as described earlier, or starting the usual GNOME desktop on the local machine by typing `startx`.

Accessing Ubuntu via Remote Desktop

A less secure but more convenient way to access your PC's desktop across a network is to use Ubuntu's Remote Desktop feature. The advantage of this method is that the desktop of the remote computer appears in a window on the remote computer's desktop, so there's no need to kill the existing X server and start a new one.

Remote Desktop uses the Virtual Network Computing (VNC) software to share the desktop. *Sharing* is the key word because, effectively, anyone who connects will take control of the main desktop. They will control the mouse and keyboard input.

However, there are a couple of important differences compared to accessing X across an ssh connection:

- Although the password is sent encrypted, the rest of the VNC data transfers aren't. Complete encryption is possible using special versions of VNC, or via an OpenSSH tunnel, but this can be difficult to set up on the Windows end of the connection.
- The remote desktop isn't blanked, so anyone standing in front of the computer will be able to see what you're doing. This could present a security/privacy risk.

If you're prepared to put up with these caveats, then allowing remote desktop access on a computer is easy. Here's the procedure:

1. Click System ► Preferences and then click Remote Desktop. Put a check alongside Allow Other Users to View Your Desktop and ensure there's a check in the box alongside Allow Other Users to Control Your Desktop, as shown in Figure 34-3. Beneath this option, you can choose whether the user can confirm each connection and whether you want to set a password. Both options add to the security of your system, although the confirmation option will mean that someone will have to be at the computer to authorize an incoming connection.



Figure 34-3. Ubuntu's Remote Desktop feature lets you share your desktop but isn't as secure as making an X server connection across SSH.

Note Ignore the information in the dialog box about how users can view your desktop, because the address it gives is wrong. Instead, you should find out your computer's IP address and use that. To do this, click Applications ► System Tools ► Network Tools. Ensure your network device is selected in the drop-down list and look for the IPv4 IP address in the list.

2. On the computer from which you want to connect to the remote desktop, click Applications ► Internet ► Terminal Server Client. In the Computer field, enter the IP address of the computer you wish to connect to. In the Protocol drop-down list, select VNC. Then click Connect.
3. You'll be prompted to enter the password, if one is applicable, and you should then see the remote computer's desktop in a window. Once you've finished, simply close the window to terminate the connection.

Connecting to Remote Windows Computers

The Terminal Server Client program allows you to connect to a variety of remote desktop server programs and, in particular, Windows XP or Windows 2000 computers via the Remote Desktop Protocol (RDP). Unfortunately, Windows XP Home, Me, 98, and 95 don't support RDP connections, which means that they aren't able to run an RDP server and allow other computers to access their desktops. However, there is a way to access the desktop of these computers remotely using some add-in software.

Connecting to Windows XP Professional, 2000, and NT

Using Terminal Server Client to access a Windows XP Professional, 2000, or NT computer is easy. Here, I use an XP Professional machine as an example, but the instructions are valid for 2000 and NT, too.

You'll find the Terminal Server Client program on the Applications ► Internet menu. Once it's running, in the Computer field, type either the IP address of the machine or its FQDN (if applicable). You don't need to type the username, password, or any other details. Click Connect, and a new window should appear in which you should see an XP login prompt. You should then log in to Windows using your username and password.

Caution If you haven't set a password for your user account on the Windows machine, you won't be able to log in. This is a quirk of the Windows XP RDP system. The solution is simple: use the User Accounts applet within the Windows Control Panel to assign yourself a password.

Of course, the XP computer will need be configured to allow incoming RDP connections. To configure it, right-click My Computer, selecting Properties, click the Remote tab, and make sure Allow Users to Connect Remotely to This Computer is checked, as shown in Figure 34-4. The Windows computer to which you want to connect may also need to be updated with the latest service packs, particularly in the case of a Windows 2000 computer.

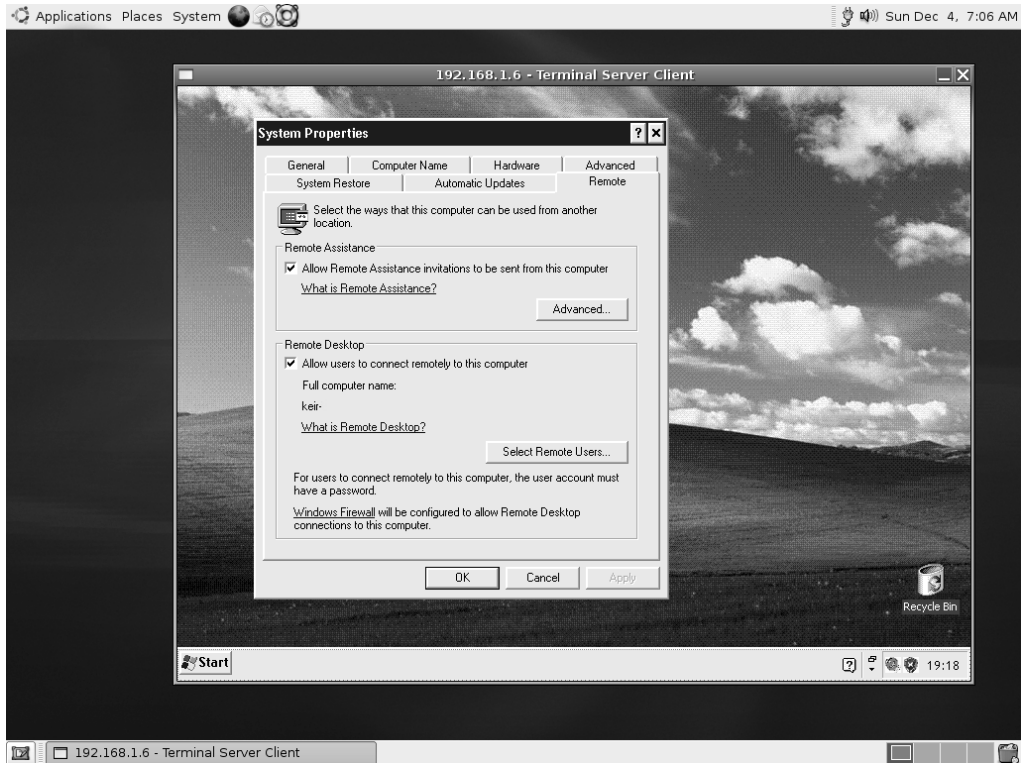


Figure 34-4. You can access remote Windows XP Professional machines using RDP and the Terminal Server Client program.

If this is the first time you've accessed the Windows computer over an RDP connection, you might be wondering why the graphics look so bad. This is because they're heavily compressed in order to transmit efficiently across networks.

Note As with Ubuntu's Remote Desktop feature, when logging in via RDP, you will be literally sharing the same desktop as anyone who might be standing in front of the computer. Because of this, for privacy and security, the desktop on the remote computer will be blanked when you log in.

Connecting to Other Windows Computers

You can download a VNC server for just about any operating system. Windows and Linux are supported, as is Macintosh OS X. In fact, a VNC server will run on any Windows computer, from 95 upwards. Once it's installed, you can then use the Terminal Server Client program within Ubuntu to connect to that computer's remote desktop.

Note In fact, any computer that's running the VNC Viewer program can access a computer running a VNC server (including the one set up by Ubuntu's Remote Desktop feature). Various VNC Viewer programs are available for Linux, Windows, Macintosh OS X, and other operating systems, including the likes of PocketPC. Just search the Web using "VNC Viewer" as a search string to find viewer programs.

Of course, you'll have the same insecurities and lack of desktop blanking that plague VNC connections to a Ubuntu desktop, as described previously. But if you're prepared to accept this, you'll be pleased to hear that setting up the VNC server on the Windows machine is easy. TightVNC, available from www.tightvnc.com, is one of the best variations of VNC around. You should download the Self-Installing Package for Windows. During installation, you'll be asked if you want to register TightVNC as a system service. Click the check box alongside this option. This will activate the VNC server every time the computer starts.

Once the program has installed, the server configuration program will appear. You should change the password by overtyping the default in the Password field.

Connecting to the remote Windows machine is also a piece of cake. On the Ubuntu system, open Terminal Server Client (Applications ► Internet) and type the remote computer's IP address into the Computer field. In the Protocol drop-down list, select VNC. Then click Connect. You'll be prompted for the remote computer's VNC server password and, once you enter this, the remote desktop will appear in a window. Figure 34-5 shows an example of connecting to a Windows 98 computer.

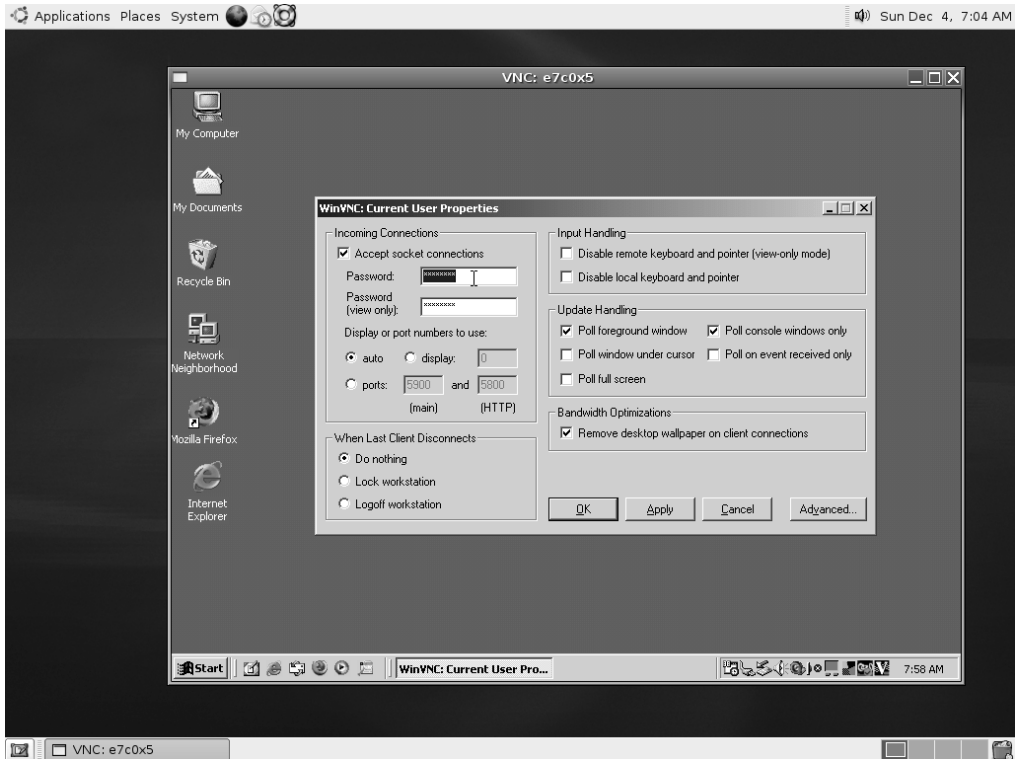


Figure 34-5. By installing a VNC server, you can connect to just about any remote operating system, including Windows 98 (as here).

Summary

In this chapter, we looked at how you can access your Ubuntu computer remotely across the Internet. We examined how you can access the computer as if you were sitting in front of it, using the `ssh` program. This allows you to start a command-line prompt and even run GUI programs on the remote computer.

In addition, we discussed how the `screen` program can be used to keep sessions alive across various logins, and how you can transfer files using the `sftp` and `scp` programs. Then we looked at how to use the Terminal Server Client tool to access the desktop of remote Windows computers.



Glossary of Linux Terms

This appendix provides brief explanations of common terms used in the Linux and Unix environments. These include technical terms, as well as conventions used within the Linux community. Due to space limitations, this glossary is somewhat selective, but still should prove a lasting reference as well as a helpful guide for those new to Linux.

Cross-referenced terms are highlighted in italics.

Symbols

.

Symbol that, in the context of file management, refers to the current directory.

..

Symbol that, in the context of file management, refers to the parent directory of that currently being browsed.

/

Symbol that, in the context of file management, refers to the *root* of the file system; also separates directories in a path listing.

~

Symbol that, in the context of file management, refers to a user's home directory.

|

Pipe symbol; used at the *command prompt* to *pipe* output from one *command* to another.

>

Symbol that, when used at the *command prompt*, indicates output should *redirect* into a file.

<

Symbol that, when used at the *command prompt*, indicates a *command* should accept input from a file (see *redirect*).

#

Symbol that, when it appears on the *command prompt*, usually indicates the user is currently logged in as *root*.

\$

Symbol that, when it appears on the *command prompt*, usually indicates the user is currently logged in as an ordinary user. (Note that some versions of *Linux/Unix* use % or > instead of \$.)

?

Wildcard character indicating that any character can be substituted in its place.

*

Wildcard character indicating that zero or more characters can appear in its place.

***nix**

Popular but unofficial way of describing the family tree that comprises *Unix* and its various clones, such as *Linux* and *Minix*.

A

administrator

Another word for either the *root* user or one who has adopted that user's powers temporarily.

AIX

IBM's *proprietary* form of *Unix* that runs on the company's *proprietary* hardware, as well as *commodity* hardware based around AMD and Intel processors. Nowadays, IBM is slowly deprecating AIX in favor of *Linux*.

alias

Method of creating a user-defined *command* that, when typed, causes another command to be run or a *string* to be expanded.

Apache

Popular *Open Source* web server software that runs on *Unix*, *Linux*, and other operating system platforms. Considered responsible in part for the rise in popularity of *Linux* in the late 1990s.

applet

Small program that, in the context of the *Ubuntu* desktop, runs as part of a larger program and offers functions that complement the main program. The *GNOME* desktop incorporates several applets in its notification area.

APT

Advanced Packaging Tool; the underlying system by which software is managed and installed on *Ubuntu* and *Debian Linux* systems. *Shell* commands beginning with *apt*, such as *apt-get*, are used to install new software from various repositories. Under *Ubuntu*, the Synaptic Package Manager program provides a *GUI* method of using APT.

archive

Any file containing a collection of smaller files, compressed or otherwise (see also *tar*).

B**BASH**

Bourne Again SHell. The most common *shell* interpreter used under *Linux* and offered as default on many Linux systems.

binary executable

Another way of referring to a program that has been compiled so that it can be used day-to-day. See also *compile*.

block device

How the *Linux kernel* communicates with a *device* that sends and receives blocks of data; usually a hard disk or removable storage device. See also *character device*.

BSD Unix

Berkeley Software Distribution Unix; form of *Unix* partially based on the original Unix *source code* but also incorporating recent developments. BSD is *open source* and free for all to use and share with practically no restrictions. There are various forms of BSD Unix, such as FreeBSD, NetBSD, and OpenBSD. BSD doesn't use the *Linux kernel*, but it runs many of the same programs. Some of the programs offered within the Linux operating system come from BSD.

bzip2

Form of file compression. Together with the older and less efficient *gzip*, it is a popular form of file compression under *Linux* and the equivalent to *Zip* compression under Windows. Files employing bzip compression are usually given a *.bz2* file extension. See also *tar*.

C

C

Programming language in which much of the *Linux kernel* is written, as were later versions of *Unix* before it. C was created by some of the same people who created Unix, and its development mirrors that of Unix.

C++

Object-oriented programming language; originally designed to be an enhancement to C, but now seen as a popular alternative.

C#

Modern programming language, which uses similar syntax to C, created by Microsoft and re-created on *Linux* via the Mono project.

character device

How Linux refers to a *device* that sends/receives data asynchronously. For various technical reasons, this typically refers to the *terminal* display. See also *block device*.

checksum

Mathematical process that can be applied to a file or other data to create a unique number relative to the contents of that file. If the file is modified, the checksum will change, usually indicating that the file in question has failed to download correctly or has been modified in some way. The most common type of checksum program used under *Linux* is md5sum.

client

Shorthand referring to a computer that connects to a *server*.

closed source

The reverse of *Open Source* in which the *source code* is not available for others to see, share, or modify. See also *proprietary*.

code

See *source code*.

command

Input typed at the *shell* that performs a specific task, usually related to administration of the system and/or the manipulation of files.

command-line prompt

See *shell*.

commodity

In the context of hardware, describes PC hardware usually based around Intel or AMD processors that can be bought off the shelf and used to create sophisticated computer systems (as opposed to buying specially designed hardware). One reason for *Linux*'s success is its ability to use commodity hardware.

community

The general term for the millions of *Linux* users worldwide, regardless of what they use Linux for or their individual backgrounds. By using Linux, you automatically become part of the community.

compile

The practice of creating a binary file from *source code*, usually achieved using the `./configure`, `make`, `make install` series of commands and scripts.

config file

Configuration file; any file that contains the list of settings for a program. Sometimes it's necessary to edit config files by hand using programs like *vi* or *Emacs*, but often the program itself will write its config file according to the settings you choose.

copyleft

The legal principle of protecting the right to share a creative work, such as a computer program, using a legally binding license. Copyleft also ensures future iterations of the work are covered in the same way.

cracker

Someone who breaks into computer systems to steal data or cause damage. The term is not necessarily linked to *Linux* or *Unix* but was created by the *community* to combat the widespread use of hacker in this sense. The word *hacker* has traditionally defined someone who merely administers, programs, and generally enjoys computers.

cron

Background *service* that schedules tasks to occur at certain times. It relies on the crontab file.

CUPS

Common Unix Printing System; set of programs that work in the background to handle printing under *Unix* and *Linux*.

curses

Library that lets software present a semigraphical interface at the *shell*, complete with menu systems and simple mouse control (if configured). The version of curses used under *Linux* and *Unix* is called ncurses.

CVS

Concurrent Versioning System; application that allows the latest version of software packages to be distributed over the Internet to developers and other interested parties.

D**daemon**

See *service*.

Debian

Voluntary organization that produces *distributions* of *Free Software* operating systems, including *Linux*. Because it is a nonprofit organization run by passionate Free Software advocates, it is considered the most ethically sound of all Linux outfits. Many *distributions*, including *Ubuntu*, use Debian as the basis for their software because of its claimed reliability and relative simplicity.

dependency

A way of referring to system files that a program requires in order to run. If the dependencies are not present during program installation, a program might refuse to install.

device

Linux shorthand describing something on your system that provides a function for the user or that the system requires in order to run. This usually refers to hardware, but it can also describe a virtual device that is created to provide access to a particular Linux function.

directory

What Windows refers to as a folder; areas on a hard disk in which files can be stored and organized.

distribution

A collection of software making up the *Linux* operating system; also known as a *distro*. The software is usually compiled by either a company or organization. A distribution is designed to be easy to install, administer, and use by virtue of it being an integrated whole. Examples include *Ubuntu*, *SUSE Linux*, *Red Hat*, and *Debian*.

distro

Shorthand for *distribution*.

documentation

Another way of describing written guides or instructions; can refer to online sources of help as well as actual printed documentation.

dpkg

Shell command that can be used to administer software under *Ubuntu* and *Debian*. However, the *APT* system, which uses `dpkg`, is the preferred method of installing software.

E**Emacs**

Seminal text editor and pseudo-*shell* beloved by *Unix* aficionados; can be used for programming tasks, simple word processing, and much more. This editor has cultural significance as one of the core pieces of software offered by *GNU Project, The*. Emacs was originally developed principally by *Stallman, Richard*. See also *vi*.

environment

Shorthand referring to a user's unique *Linux* configuration, such as variables that tell the *shell* where programs are located.

F**FAT32**

File Allocation Table 32-bits; file system offered by Windows 98, Me, 2000, and XP. *Linux* can both read and write to FAT32 file systems. See also *NTFS* and *VFAT*.

Firefox

Web browser program used under *Ubuntu* and produced by the *Mozilla Foundation*.

FLOSS

Free, Libre, or Open Source Software; used within the *community* to describe all software or technology that, broadly speaking, adheres to the ethical approach of *Open Source* software and/or *Free Software*, as well as its legal guidelines.

FOSS

Free or Open Source Software; alternative term for *FLOSS*.

free

When used to describe software or associated areas of technology, “free” indicates that the project abides by the ethical (if not legal) guidelines laid down by *GNU Project, The*. It doesn't indicate that the software is free in a monetary sense; its meaning is quite different from “freeware.”

Free Software

Software in which the *source code*—the original listing created by the programmer—is available for all to see, share, study, and adapt to their own needs. This differs from the concept of *open source* because the right of others to further modify the code is guaranteed via the *GNU Public License* (GPL) software license (or a compatible license). For various reasons, Free Software sometimes does not include the source code (although the software can still be legally decompiled), but this is rare.

G

gcc

GNU Compiler Collection; programs used when creating *binary executable* files from *source code*.

GID

Group ID; numbering system used by the operating system to refer to a *group*.

GIMP

GNU Image Manipulation Program; high-powered image-editing program that runs under *Linux*, *Unix*, Windows and other operating systems. Often preceded by the definite article: “The GIMP.”

GNOME

GNU Network Object Model Environment; a *GUI*-based desktop environment used by *Ubuntu*, as well as several other *distributions*. It uses the *GTK+* *libraries*. See also *KDE*.

GNU

GNU’s Not Unix; see *GNU Project, The*.

GNU/Linux

Another name for the operating system referred to as *Linux*. The name GNU/Linux gives credit to the vast quantity of *GNU Project, The* software that is added to the *Linux kernel* within a *distro* to make a complete operating system. As such, GNU/Linux is the preferred term of many *Free Software* advocates.

GNU Project, The

Organization created by *Stallman, Richard* in order to further the aims of *Free Software* and create the body of software that makes up the *GNU* operating system.

GNU Public License

Software license principally created by *Stallman, Richard* in order to protect software *source code* against *proprietary* interests and ensure that it will always be shared. It does this by insisting that any source code covered by the GNU Public License (GPL) must remain licensed under the GPL, even after it has been modified or added to by others. The *Linux kernel*, as well as much of the software that runs on it, uses the GPL.

GPL

See *GNU Public License*.

grep

Global Regular Expression Print; powerful *shell command* that lets you search a file or other form of input using *regular expressions*. Because of the ubiquity of the grep program, many *Linux* and *Unix* users refer to searching as “grepping.” To “grep a file” is to search through it for a *string*.

group

Collection of users under one heading (group name) to facilitate system administration.

GRUB

GRand Unified Bootloader; boot manager program that offers a menu from which you can choose which operating system you wish to boot. It’s needed to load the *kernel* program and thereby initiate the *Linux* boot procedure.

GUI

Graphical user interface; describes the software that provides a graphical system to display data and let you control your PC (usually via a mouse).

guru

One who is experienced and knowledgeable about *Linux/Unix* and is willing to share his or her knowledge with others. In a perfect world, every *newbie* would have his or her own guru.

gzip

One of the two preferred forms of file compression used under *Linux*. Files employing gzip compression usually have a *.gz* file extension. See also *bzip2*.

H

hack

Ingenious and/or extremely efficient solution to a problem, particularly within the programming world.

hacker

Term used within the *community* to describe anyone who enjoys computers and possesses some skill therein, either in a professional capacity or as a hobby. This term is distinct from connotations of maliciously breaking into computers propagated by the media. See also *cracker*.

host

Shorthand referring to any computer that acts as a *server* to another computer. See also *client*.

HP-UX

Hewlett-Packard's *proprietary* form of *Unix* designed to work on its own hardware platform.

Hurd

Kernel being developed by *GNU Project, The*. It's not associated with the *Linux* kernel in any way.

I**info**

Source of *documentation* accessible from the *shell*; an alternative to the more established *man page* system. Also known as *Texinfo*.

init

The program that is automatically run after the *kernel* has finished loading, and therefore early in the boot procedure. It's responsible for effectively starting the operating system.

init.d

Collection of startup *scripts* that make up the components of a *run level*. Under *Ubuntu*, these are found at `/etc/init.d/`. *Symbolic links* to selected *init.d* scripts are contained in folders within `/etc/init.d` that are named after *run level* numbers, such as `rc0.d`, `rc1.d`, `rc2.d`, and so on.

initrd

Initial RAM disk; system used by the *Linux kernel* to load *modules* that are essential for the kernel to be able to boot, such as disk controllers.

inode

Part of the usually invisible file system structure that describes a file, such as its ownership permissions or file size.

ipchains

Now deprecated component of version 2.2 of the *Linux kernel* that allows the creation of network security setups, such as firewalls or port-forwarding arrangements. Note that some *distros* still prefer to use *ipchains*. See also *iptables*.

iptables

Component of versions 2.4 and 2.6 of the *Linux kernel* that allows powerful network security setups. Chiefly used in the creation of firewalls, but can be used for more elementary arrangements such as network address translation (NAT) routers.

J

job

How the *BASH shell* refers to a running program in order to facilitate administration by the user.

journaling

File system technology in which integrity is maintained via the logging of disk writes.

K

KDE

K desktop environment; *GUI* and set of additional programs used on various *distros*, such as *Mandriva* and a variation of *Ubuntu* called *Kubuntu*.

kernel

Essential but ordinarily invisible set of programs that run the computer's hardware and provide a platform on which to run software. In the *Linux* operating system, the kernel is also called *Linux*, after its creator, *Torvalds*, *Linus*.

kernel panic

Error message that appears when the *kernel* program in *Linux* cannot continue to work. In other words, a polite way of indicating a crash or, more often, a problem arising from user misconfiguration. This is most often seen when booting up after making incorrect changes to the system.

kludge

Community slang describing an inelegant way of making something work, usually not in a way that is generally accepted as being correct. Pronounced "kloodge."

Kubuntu

Version of *Ubuntu* that substitutes the *GNOME* desktop environment for *KDE*. At the time of writing, it is less mature than the standard version of *Ubuntu*.

L

LAMP

Acronym describing a series of programs that work together to provide a complete *Linux*-based web-hosting environment. Stands for *Linux*, *Apache*, *MySQL*, and Perl, PHP or Python (the last three in the list are scripting languages; see *script*).

LGPL

Lesser GPL; version of the *GNU Public License* (GPL) in which some use restrictions are slackened at the expense of various freedoms laid down by the main GPL. The LGPL is mostly used for *library* files.

library

General term referring to code that programs need to run and that, once in memory, is frequently accessed by many programs (leading to the phrase “shared library”). The most common and vital library is *glibc* (GNU C Library), created by *GNU Project*, *The* and the fundamental building block without which *Linux* could not operate. *GNOME* relies on the *GTK+* libraries, among others.

link

File system method of assigning additional filenames to a file; also known as a “hard link.” See also *symbolic link*.

Linux

You mean you don’t know by now? *Linux* is what this book is all about. It is a *kernel* program created by *Torvalds*, *Linus* in 1991 to provide an inexpensive operating system for his computer, along with other components. These days, *Linux* is used to describe the entire operating system discussed in this book, although many argue (perhaps quite rightly) that this is inaccurate, and use the term *GNU/Linux* instead.

local

Shorthand referring to the user’s PC or a device directly attached to it (as opposed to *remote*).

localhost

(1) Network name used internally by *Linux* and software to refer to the *local* computer, distinct from the network.

(2) Default name given to a Linux-based PC when no other name is defined during installation. However, under Ubuntu, the name *ubuntu* is assigned if no name is defined by the user.

M

man page

Documentation accessible from the *shell* that describes a *command* and how it should be used.

Minix

Operating system that is a rough clone of *Unix*, created by Professor Andrew Tanenbaum. It was the inspiration for *Linux*.

module

Program code that can be inserted or removed from the *kernel* in order to support particular pieces of hardware or provide certain kernel functions. Drivers under Windows perform the same function.

mount

To add a file system so that it is integrated (and therefore accessible) within the main file system; applies to external file systems, such as those available across networks, as well as those on the *local* PC, such as the hard disk or CD/DVD-ROMs.

Mozilla Foundation

Organization founded by Netscape to create *open source* Internet software, such as web browsers and e-mail clients; originally based on the Netscape *source code*. At the time of writing, it produces the *Firefox* and Camino web browsers, the Thunderbird e-mail and Usenet client, the Bugzilla bug-tracking software, as well as other programs.

MySQL

Popular and powerful *open source* database application. See also *LAMP*.

N

newbie

Term used to describe anyone who is new to *Linux* and therefore still learning the basics. It's not a derogatory term! See also *guru*.

NFS

Network File System; reliable and established method of sharing files, printers, and other resources across a network of *Unix*-based operating systems. See also *Samba*.

NTFS

NT File System; file system offered by Windows NT, 2000, and XP. It can be read by *Linux*, but usually writing is prohibited because it is considered unsafe. See also *FAT32*.

O**OpenOffice.org**

Open Source office suite project created with the continuing input of Sun Microsystems and based on code Sun contributed to the Open Source *community*. Its commercial release is in the form of Star Office (although Star Office has several *proprietary* components added).

open source

(1) Method and philosophy of developing software whereby the *source code*—the original listing created by the programmer—is available for all to see. Note that open source is not the same as *Free Software*; describing software as open source doesn't imply that the code can be shared or used by others (although this is often the case).

(2) A community of users or any project that adheres to open-source values and/or practices.

P**partition**

Subdivision of a hard disk into which a file system can be installed.

PID

Process ID; the numbering system used to refer to a *process*.

pipe

Method of passing the output from one *command* to another for further processing. Piping is achieved within the *shell* by typing the | symbol.

POSIX

Portable Operating System Interface; various technical standards that define how *Unix*-like operating systems should operate and to which the *Linux* operating system attempts to adhere.

PPP

Point-to-Point Protocol; networking technology that allows data transfer across serial connections like telephone lines. In other words, it's the technology that lets you connect to your Internet service provider using a modem.

process

The way the system refers to the individual programs (or components of programs) running in memory.

proprietary

Effectively, software for which a software license must be acquired, usually for a fee. This usually means the *source code* is kept secret, but it can also indicate that the source code is available to view but not to incorporate into your own projects or share with others.

R

Red Hat

Well-known company that produces distributions of *Linux*.

redirect

To send the output of a *command* into a particular file. This also works the other way around: the contents of a particular file can be directed into a command. Redirection is achieved within the *shell* using the left and right angle brackets (< and >), respectively.

regex

See *regular expression*.

regular expression

Powerful and complex method of describing a search *string*, usually when searching with tools such as *grep* (although regular expressions are also used when programming). Regular expressions use various symbols as substitutes for characters or to indicate patterns.

remote

Indicates a computer or *service* that is available across a network, including but not limited to computers on the Internet (as opposed to *local*).

root

- (1) The bottom of the *Linux* file system directory structure, usually indicated by a forward slash (/).
- (2) The user on some versions of *Unix* or *Linux* who has control over all aspects of hardware, software, and the file system.
- (3) Used to describe a user who temporarily takes on the powers of the root user (via the `sudo` command, for example).

RPM

Red Hat Package Manager; system used to install and administer programs under *Red Hat*, *SUSE Linux* and some other *distributions*. See also *APT*.

RTFM

Read the freaking manual / *man page*; exclamation frequently used online when a *newbie* asks for help without having undertaken basic research.

run level

Describes the current operational mode of *Linux* (typically, the *services* that are running). Run level 1 is single-user mode (a stripped-down system with minimal running services); run levels 2 through 5 provide a *GUI*; run level 6 is reboot mode (switching to it will cause the computer to terminate its processes and then reboot); run level 0 is shutdown (switching to it will cause the PC to shutdown).

S**Samba**

Program that re-creates under *Unix* or *Linux* the Microsoft *SMB*-based system of sharing files, printers, and other computer resources across a network. It allows Linux to become a file or printer server for Linux and Windows computers, and also allows a Linux client to access a Windows-based server.

scalable

Term describing the ability of a single computer program to meet diverse needs, regardless of the scale of the potential uses. The *Linux kernel* is described as being scalable because it can run supercomputers, as well as handheld computers and home entertainment devices.

script

Form of computer program consisting of a series of *commands* in a text file. Most *shells* allow some form of scripting, and entire programming languages such as Perl are based around scripts. In the context of the Linux operating system, shell scripts are usually created to perform trivial tasks or ones that frequently interact with the user. Shell scripts have the advantage that they can be frequently and easily modified. The *Linux* boot process relies on several complex scripts to configure essential system functions such as networking and the *GUI*. See also *init*.

server

- (1) Type of computer designed to share data with other computers over a network.
- (2) Software that runs on a computer and is designed to share data with other programs on the same PC or with other PCs across a network.

service

Background program that provides vital functions for the day-to-day running of *Linux*; also known as a *daemon*. Services are usually started when the computer boots up and as such are constituent parts of a *run level*.

shell

Broadly speaking, any program that creates an operating environment in which you can control your computer. The *GNOME* desktop can be seen as a shell, for example. However, it's more commonly understood within *Unix* and *Linux* circles as a program that lets you control the system using *commands* entered at the keyboard. In this context, the most common type of shell in use on Linux is *BASH*.

Shuttleworth, Mark

Entrepreneurial South African businessman who, as a long-term *Debian hacker*, devised and financially supports *Ubuntu* via his company, Canonical Ltd.

SMB

Server Message Block; network technology for sharing files, printers, and other resources. See also *Samba*.

Solaris

Form of *Unix* sold by Sun Microsystems; runs on *proprietary* hardware systems as well as on *commodity* systems based on Intel and AMD processors.

source code

The original program listing created by a programmer. Most programs that you download are precompiled—already turned into *binary executables* ready for general use—unless you specifically choose to download and *compile* the source code of a program yourself.

SSH

Secure SHell; program that lets you access a *Linux/Unix* computer across the Internet. SSH encrypts data sent and received across the *link*.

SSL

Secure Sockets Layer; form of network data transfer designed to encrypt information for security purposes. It's used online for certain web sites and also within *Linux* for certain types of secure data exchange.

Stallman, Richard M.

Legendary *hacker* who founded *GNU Project, The* and created the concept of *copyleft*, as well as the software license that incorporates it: the *GNU Public License (GPL)*. See also *Torvalds, Linus*.

standard error

Linux and *Unix* shorthand for the error output provided by a *command*.

standard input

Linux and *Unix* shorthand for the *device* usually used to provide input to the *shell*. For the majority of desktop PC users, this refers to the keyboard.

standard output

Linux and *Unix* shorthand for the *device* usually used to display output from a *command*. For the majority of desktop PC users, this refers to the screen.

string

A word, phrase, or sentence consisting of letters, numbers, or other characters that is used within a program and is often supplied by the user.

sudo

Program that runs under *Unix* and *Linux* by which ordinary users are temporarily afforded *administrator* rights. Ubuntu relies on *sudo* as the exclusive way for users to administer the system.

SVG

Scalable Vector Graphics; vector graphics technology. SVG is actually an XML markup language designed to create 2D graphics, increasingly used for *Linux* desktop icons and web graphics.

swap

Area of the hard disk that the *Linux kernel* uses as a temporary memory storage area. Desktop or *server* Linux differs from Windows in that it usually requires a separate hard disk *partition* in which to store the swap file.

symbolic link

Type of file akin to a Windows shortcut. Accessing a symbolic link file routes the user to an actual file. See also *link*.

sysadmin

Systems administrator; a way of describing the person employed within a company to oversee the computer systems. In such an environment, the sysadmin usually is the *root* user of the various computers.

System V

Variant of *Unix* used as a foundation for modern forms of *proprietary* Unix.

T**tainted**

Describes a *kernel* that is using *proprietary modules* in addition to *Open Source* modules. Can also refer to insecure software.

tar

Tape Archive; software able to combine several files into one larger file in order to back them up to a tape drive or simply transfer them across the Internet. Such files are usually indicated by a `.tar` file extension. Note that a tar file isn't necessarily compressed; the *bzip2* and *gzip* utilities must be used if this is desired.

TCP/IP

Transmission Control Protocol/Internet Protocol; standard protocol stack used by most modern operating systems to control and communicate across networks and also across the Internet (as opposed to NetBEUI, commonly available on older versions of Windows, and IPX/SPX, used on Novell's NetWare operating system).

terminal

Another word for *shell*.

TeX

Method and set of programs for typesetting complex documents. Invented prior to word processors and desktop publishing software, and now considered a specialized tool for laying out scientific texts. An updated version of the program called LaTeX is also available.

Torvalds, Linus

Finnish programmer who, in 1991, created the initial versions of the *Linux kernel*. Since then, he has taken advantage of an international network of volunteers and staff employed by various companies who help produce the kernel. Torvalds himself contributes and oversees the efforts.

tty

TEleTYpewriter; shorthand referring to underlying *Linux* virtual *devices* that allow programs and users to access the *kernel* and thereby run programs.

Tux

The name of the penguin character that is the *Linux* mascot. The original Tux graphic was drawn by Larry Ewing.

U

Ubuntu

Linux distribution with several unique characteristics. Ubuntu is designed primarily for desktop use, although several server versions are available. It is intended for use by individuals in any location in the world, so it has strong multiple language support. It's run by the *Ubuntu Foundation*, which is financially backed by *Shuttleworth, Mark*. Each release is guaranteed to be supported for 18 months (5 years in the case of server releases).

Ubuntu Foundation

Organization set up by *Shuttleworth, Mark* and his company, Canonical, Ltd., to administer and provide support for the *Ubuntu* distribution of *Linux*.

UID

User ID; numbering system used by the operating system to refer to a *user*.

Unix

Seminal operating system created as a research project in 1969 by Kenneth Thompson and Dennis Ritchie at Bell Labs (later AT&T). Because it was initially possible to purchase the *source code* for a fee, subsequent revisions were enhanced by a variety of organizations and went on to run many mainframe and minicomputer systems throughout the 1980s, 1990s, and up to the present. Nowadays, Unix is fragmented and exists in a variety of different versions. Perhaps most popular is its *Open Source* rendition, *BSD Unix*, which has seen many developments since the source code was first released. This means that BSD Unix no longer exists but has instead diversified into a number of separate projects. *Proprietary* versions are also available, including *Solaris*, *HP-UX*, and *AIX*.

user

The way the operating system refers to anyone who accesses its resources. A user must first have a user account set up, effectively giving that user his or her own private space on the system. In addition to actual human users, an average *Linux* system has many other user accounts created to let programs and *services* go about their business. These are usually not seen by human users.

V**variable**

A changeable value that stores a certain data type (such as a number, date, or *string*), remembering it for future reference by the system or *script* it is defined by. Variables defined by and for the *Linux kernel* are vital to it.

verbose

Command option that will cause it to return more detailed output (or, in some cases, to return actual output if the command is otherwise “quiet”); usually specified by adding the `-v` command option.

VFAT

Virtual File Allocation Table; technical name of Microsoft’s FAT file system offered under Windows and also on removable storage devices such as flash memory cards.

vi

Arcane text editor and pseudo-shell beloved by *Unix* aficionados that can be used for creation of text files or for creating programs. Traditionally, Unix users either love or hate vi; some prefer *Emacs*.

W**Wine**

Short for Wine Is Not an Emulator; software that re-creates the Windows Application Programming Interface (API) layer within *Linux* and lets users run Windows programs.

workspace

X terminology referring to a *GUI* desktop.

X

X

Short for X Window; software that controls the display and input devices, thereby providing a software foundation on top of which desktop managers like *GNOME* are able to run.

X11

Version 11 of the X software, currently in use on most desktop *Linux* systems.

XFree86 Project

Organization that creates X software. At one time, every *distribution* of *Linux* used XFree86 software, but most now use similar software from the *X.org* organization.

xinetd

The *service* responsible for starting various network servers on the computer.

XMMS

Audio player program.

X.org

Organization that produces the X Window software and, in particular, a set of programs called X11. X11 is used on most modern distributions of *Linux*. It is backed by a number of *Unix* and Linux industry leaders.

xterm

Simple program that allows you to run a shell under X. This program has the advantage of being available on most *Linux* systems that offer a *GUI*.



BASH Command Index

This appendix provides a whistle-stop tour of commands that can be used at the BASH shell. This is a highly selective listing, intended to provide a guide to commands that see day-to-day use on average desktop systems. In a similar fashion, although some command options are listed, they're strictly limited to those that receive regular deployment.

The descriptions of each command are deliberately simple. Note that the quantity of space a command is given is not an indication of its importance or usefulness. To this end, commands in the list with an asterisk after their name offer far more than hinted at by its brief description hints. In such cases, I strongly advise that you refer to the command's man page for more information.

Various conventions are used in the list:

- You should substitute your own details wherever italicized words appear.
- Commands that can and might be run by ordinary users are preceded with a dollar sign (\$).
- Commands that require root privileges (the use of `sudo`) are preceded with a hash symbol (#).

Commands that present dangers to the system through misuse are clearly marked. Such commands should not be used without research into the command's usage and function.

Command	Description	Typical Command Options	Examples of Use
\$ alias	Create or display command aliases		alias list=ls
\$ alsamixer	Alter audio volume levels		alsamixer
\$ apropos	Search man pages for specified words/phrases		apropos "word or phrase"

Command	Description	Typical Command Options	Examples of Use
\$ apt-cache	Search, query, and otherwise manipulate the APT database cache (see apt-get)	<p>search: Search for specified package (regexes may be used; see Chapter 15)</p> <p>showpkg: Show information about specified package</p> <p>depends: Show package dependencies of specified package, and show other packages that can meet that dependency</p>	<p>apt-cache search <i>packagename</i></p>
# apt-get	Multifunction tool use to install, remove, and otherwise administer software packages, according to the APT database	<p>install: search for and install specified package from repositories (as specified in /etc/apt/sources.list)</p> <p>update: Update or build package database by contacting package repositories</p> <p>upgrade: Attempt to upgrade all current installed packages with new versions</p> <p>dist-upgrade: Attempt to upgrade all currently installed packages, automatically and aggressively resolving package conflicts; often used to upgrade entire distro to new version</p> <p>remove: Opposite of install; removes packages</p> <p>clean: Remove any old package installation files that are stored on hard disk</p> <p>-f: Attempt to fix broken package dependencies (used with install or remove)</p> <p>--force-yes: Override any errors and thereby bypass apt-get's protective measures. Dangerous option—use with care!</p>	<p>apt-get install <i>packagename</i></p>

Command	Description	Typical Command Options	Examples of Use
\$ bzip2	Compress specified file (replaces original file with compressed file and gives it .bz2 file extension)	-d: Decompress specified file -k: Don't delete original file -t: Test; do a dry run without writing any data	bzip2 <i>myfile</i>
\$ bzip2recover	Attempt recovery of specified damaged .bz2 file		bzip2recover <i>myfile.tar.bz2</i>
\$ cal	Display calendar for current month (or specified month/year)		cal 4 2005
\$ cat	Display a file on screen or combine and display two files together		cat <i>myfile</i>
\$ cd	Change to specified directory		cd <i>/usr/bin</i>
\$ cdparanoia *	Convert CD audio tracks to hard disk files	-B: Batch mode; convert all tracks to individual files -S: Set CD read speed (2, 4, 8, 12, and so on; values relate to CD-drive spin speed; used to avoid read errors)	cdparanoia -S 8 -B
# cdrecord *	Burn audio or CD-R/RW data discs (the latter usually based on an ISO image; see mkisofs)	-dev=: Specify the drive's device number (can be discovered by running cdrecord with the scanbus option) --scanbus: Scan to see which CD-R/RW drives are present and return device numbers -speed=: Specify the write speed (2, 4, 6, 8, and so on) -v: Verbose output; obligatory for feedback on cdrecord's progress	cdrecord <i>dev=0,0,0 -speed=16 -v myfile.iso</i>
# cfdisk *	DANGEROUS! Menu-based disk-partitioning program		cfdisk <i>/dev/hda</i>
# chgrp	Change group ownership of a file/directory	-R: Recursive; apply changes to subdirectories	chgroup <i>mygroup myfile</i>

Command	Description	Typical Command Options	Examples of Use
\$ chmod	Change permissions of a file/directory (where a=all, u=user, g=group; and r=read, w=write, x=executable)	-R: Recursive; apply to subdirectories --reference=: Copy permissions from specified file	chmod a+rw <i>myfile</i>
\$ chown	Change file ownership to specified username	-R: Recursive; apply to subdirectories	chown <i>username myfile1</i>
# chroot	Change the root of the file system to the specified path		chroot <i>/home/mydirectory</i>
# chvt	Switch to the specified virtual terminal (equivalent of holding down Ctrl+Alt and pressing F1–F6)		chvt 3
\$ clear	Clears terminal screen and places cursor at top		clear
\$ cp	Copy files	-r: Recursive; copy subdirectories and the files therein -s: Create symbolic link instead of copying	cp <i>myfile1 directory/</i>
\$ crontab	Edit or display the user's crontab file (scheduled tasks)	-e: Edit the crontab file (create/amend) -l: List crontab entries -r: Delete the crontab file -u: Specify a user and edit their crontab file	crontab -e
\$ date	Display the date and time		date
\$ df	Display free disk space within file system	-h: Human readable; display sizes in KB, MB, GB, and TB, as appropriate -l: Restrict to local file systems, as opposed to network mounts	df -h
\$ diff	Display differences between specified files	-a: Consider all files text files (don't halt when asked to compare binary files) -i: ignore lowercase and uppercase differences	diff <i>myfile1 myfile2</i>
\$ diff3	Display differences between three specified files		diff3 <i>myfile1 myfile2 myfile3</i>
\$ dig	Look up IP address of specified domain		dig <i>mysite.com</i>

Command	Description	Typical Command Options	Examples of Use
\$ dmesg	Display kernel message log		dmesg
# dosfsck *	Check and repair MS-DOS–based file hard disk partition (see also fsck)	-a: Repair without asking user for confirmation -r: Repair file system asking user for confirmation when two or more repair methods are possible -v: Verbose; display more information	dosfsck -rv /dev/hda4
# dpkg	Install, remove, and otherwise administer local installation packages (on your hard disk); see also apt-get	-i: Install specified package -r: Remove (uninstall) specified package -I: Show info about specified package --ignore-depends= <i>packagename.deb</i> : Don't halt on package dependency issues (dangerous!)	dpkg -i <i>packagename.deb</i>
# dpkg-reconfigure	Reconfigure an already installed package		dpkg-reconfigure <i>packagename</i>
\$ du	Show sizes of files and folders in kilobytes	-h: Human readable; produce output in MB, GB, and TB -s: Summary; display totals only for directories rather than for individual files	du -h /home/myuser
\$ eject	Eject a removable storage disk	-t: Close an already open tray	eject /media/dvd-rom
\$ ex *	Start a simple text-editor program used principally within shell scripts		ex <i>myfile.txt</i>
\$ exit	Log out of shell (end session)		exit
\$ fdformat	Low-level format a floppy disk (this won't create a file system; see also <i>mkfs</i>)		fdformat /dev/fd0
# fdisk *	DANGEROUS! Hard-disk partitioning program	-l: List partition table	fdisk /dev/hda
\$ fg	Brings job running in background to foreground		fg 1
\$ file	Display information about specified file, such as its type		file <i>myfile</i>

Command	Description	Typical Command Options	Examples of Use
\$ find *	Find files by searching directories (starting in current directory)	-maxdepth: Specify the number of subdirectories levels to delve into, starting from 1 (current directory) -name: Specify name of file to search for -type: Specify file types to be returned; -type d returns directories and -type f returns only files	find -name "myfile"
\$ free	Display information about memory usage	-m: Show figures in MB -t: Total the columns at bottom of table	free -m
# fsck *	Check file system for errors (usually run from rescue disc)		fsck /dev/hda1
\$ ftp *	FTP program for uploading/downloading to remote sites		ftp ftp.mysite.com
\$ fuser	Show which processes are using a particular file or file system	-v: Verbose; detailed output	fuser -v myfile
\$ grep *	Search specified file for specified text string (or word)	-i: Ignore uppercase and lowercase differences -r: Recursive; delve into subdirectories (if applicable) -s: Suppress error messages about inaccessible files and other problems	grep "phrase I want to find" myfile.txt
# groupadd	Create new group		groupadd mygroup
# groupdel	Delete specified group		groupdel mygroup
\$ groups	Display groups the specified user belongs to		groups myuser
\$ gzip	Compress files and replace original file with compressed version	-d: Decompress specified file -v: Verbose; display degree of compression	gzip myfile
# halt	Initiate shutdown procedure, ending all processes and unmounting all disks	-p: Power off system at end of shutdown procedure	halt -p

Command	Description	Typical Command Options	Examples of Use
# hdparm *	DANGEROUS! Tweak or view hard disk settings		hdparm /dev/hda
\$ head	Print topmost lines of text files (default is first 10 lines)	-n: Specify number of lines (such as -n 5)	head <i>myfile.txt</i>
\$ help	Display list of common BASH commands		help
\$ history	Display history file (a list of recently used commands)		history
\$ host	Query DNS server based on specified domain name or IP address	-d: Verbose; return more information -r: Force name server to return its cached information rather than query other authoritative servers	host 82.211.81.166
\$ hostname	Display localhost-style name of computer		hostname
\$ id	Display username and group info of specified user (or current user if none specified)		id <i>myuser</i>
# ifconfig *	Display or configure settings of a network interface (assign an IP address, subnet mask, and activate/deactivate it)	down: Disable interface (used at end of command chain) netmask: Specify a subnet mask up: Enable interface (used at end of command chain)	ifconfig <i>eth0</i> 192.168.0.10 netmask 255.255.0.0 up
\$ info *	Display info page for specified command		info <i>command</i>
# init	Change current run level		init 1
\$ jobs	Display list of jobs running in background		jobs
\$ kill	Kill specified process		kill 1433
\$ killall	Kill process(es) that have specified name(s)	-i: Confirm before killing process -v: Verbose; report if and when successful	killall <i>processnumber</i>
\$ last	Display details of recent logins, reboots, and shutdowns		last
\$ ldd	Display system files (libraries) required by specified program		ldd /usr/bin/ <i>program</i>

Command	Description	Typical Command Options	Examples of Use
\$ less	Interactively scroll through a text file	-q: Quiet; disable beeps when end of file is reached or other error encountered -i: Ignore case; make all searches case-insensitive unless uppercase letters are used	less <i>myfile.txt</i>
\$ ln	Create links to specified files, such as symbolic links	-s: Create symbolic link (default is hard link)	ln -s <i>myfile1 myfile2</i>
\$ lpr	Print file (send it to the printer spool/queue)	-V: Verbose; print information about progress of print job	lpr <i>myfile.txt</i>
\$ lpstat	Display print queue		lpstat
\$ ls	List directory	-a: List all files, including hidden files -d: List only directory names rather than their contents -h: Human readable; print figures in KB, MB, GB, and TB -l: Long list; include all details, such as file permissions -m: Show as comma-separated list	ls -h <i>mydirectory</i>
# lsmod	Display currently loaded kernel modules		lsmod
\$ lsof	Display any files currently in use	-u: Limit results to files used by specified user	lsof -u <i>username</i>
\$ man	Display specified command's manual		man <i>command</i>
\$ md5sum	Display MD5 checksum (normally used to confirm a file's integrity after download)		md5sum <i>myfile</i>
# mkfs *	DANGEROUS! Create specified file system on specified device (such as a floppy disk)	-t: Specify type of file system	mkfs -t <i>vfat /dev/fd0</i>

Command	Description	Typical Command Options	Examples of Use
\$ mkisofs *	Create ISO image file from specified directory (usually for burning to disc with cdrrecord)	-o: Options; this must appear after command to indicate that command options follow -apple: Use Mac OS extensions to make disc readable on Apple computers -f: Follow symbolic links and source actual files -J: Use Joliet extensions (make ISO compatible with Windows) -R: Use Rock Ridge extensions (preferred Linux CD-ROM file system) -v: Verbose; display more information (-vv for even more info)	mkisofs -o <i>isoimage.iso</i> -R -J -v <i>mydirectory</i>
# modinfo	Display information about kernel module		modinfo <i>modulename</i>
# modprobe	Insert specified module into the kernel, as well as any others it relies on	-k: Set module's autoclean flag so it will be removed from memory after inactivity -r: Remove specified module as well as any it relies on to operate	modprobe <i>modulename</i>
\$ more	Interactively scroll through text file (similar to less)		more <i>myfile.txt</i>
# mount *	Mount specified file system at specified location	-o: Specify command options, such as <i>rw</i> to allow read/write access; various types of file systems have unique commands	mount <i>/dev/hda4 /mnt</i>
\$ mv	Move (or rename) specified files and/or directories	-b: Back up files before moving -v: Display details of actions carried out	mv <i>myfile mydirectory/</i>
\$ netstat *	Show current network connections		netstat -a

Command	Description	Typical Command Options	Examples of Use
\$ nice	Run specified command with specified priority	-n: Specify priority, ranging from the highest priority of -20, to 19, which is the lowest priority	nice -n 19
\$ nohup	Run specified command and continue to run it, even if user logs out		nohup <i>command</i>
\$ passwd	Change user's password		passwd
\$ ping	Check network connectivity between local machine and specified address	-w: Exit after specified number of seconds (such as -w 5)	ping <i>mydomain.com</i>
\$ printenv	Display all environment variables for current user		printenv
\$ ps *	Display currently running processes	a: List all processes (note that command options don't require preceding dash) f: Display ownership of processes using tree-style graphics u: Limit results to processes running for and started by current user x: Include processes in results not started by user but running with the user ID	ps aux
\$ pwd	Display current directory		pwd
# reboot	Reboot computer		reboot
\$ renice	Change a process's priority while it's running (see nice)		renice 19 10704
\$ rm	Delete single or multiple files and/or directories	-r: Recursive; delete specified directories and any subdirectories -f: Force; don't prompt for confirmation before deleting (use with care!)	rm -rf <i>mydirectory</i>
# rmmmod	Delete module from kernel		rmmmod <i>modulename</i>
# route *	Add and create (or view) entries in routing table (see ifconfig)		route add default gw 192.168.1.1
\$ runlevel	Display current run level		runlevel

Command	Description	Typical Command Options	Examples of Use
\$ screen *	Program that runs pseudo shell that is kept alive regardless of current user login	-ls: Display list of currently running screen sessions -R: Reattach to already running screen session or start new one if none available	screen
\$ sftp *	Secure Shell FTP; like FTP but running over an ssh connection (see ssh)		sftp username@192.168.1.14
\$ shred	Overwrite data in a file with gibberish, thereby making it irrecoverable	-u: Delete file in addition to overwriting -v: Verbose; show details of procedure -f: Force permissions to allow writing if necessary	shred -fv myfile
\$ sleep	Pause input for the specified period of time (where s=seconds, m=minutes, h=hours, d=days)		sleep 10m
\$ smbclient *	FTP-style program with which you can log in to a SMB (Windows)-based file share		smbclient //192.168.1.1/
\$ sort	Sort entries in the specified text file (default is ASCII sort)		sort myfile.txt -o sorted.txt
\$ ssh *	Log in to remote computer using secure shell		ssh username@192.168.1.15
\$ startx	Start GUI session (if GUI isn't already running)		startx
\$ su	Temporarily log in as specified user; log in as root if no user specified (provided root account is activated)	-: Adopt user's environment variables, such as \$PATH	su
\$ sudo	Execute specified command with root privileges		sudo command
\$ tac	Display specified text file but in reverse (from last to first line)		tac myfile.txt
\$ tail	Display final lines of specified text file	-n: Specify number of lines to display (such as -n4)	tail myfile.txt

Command	Description	Typical Command Options	Examples of Use
\$ tar *	Combine specified files and/or directories into one larger file, or extract from such a file	-c: Create new archive -j: Use bzip2 in order to compress (or decompress) files -f: Specifies filename (must be last in chain of command options) -r: Add files to existing archive -x: Extract files from existing archive -z: Use gzip to compress (or decompress) files	tar -zcf myfile.tar.gz mydirectory
\$ tee	Display piped output and also save it to specified file		ls -lh tee listing.txt
\$ top *	Program that both displays and lets the user manipulate processes		top
\$ touch	Give specified file current time and date stamp; if it doesn't exist, create a zero-byte file with that name		touch myfile
\$ tracepath	Discover and display network path to another host		tracepath 192.168.1.20
\$ umask	Set default permissions assigned to newly created files		umask u=rwx,g=r,o=
# umount	Unmount a file system		umount /media/cdrom
# useradd	Add new user	-m: Create home directory for user	useradd -m username
# userdel	Delete all mention of user in system configuration files (effectively deleting the user, although files owned by the user might remain)	-r: Remove user's home directory	userdel -r username
\$ unalias	Remove specified alias	-a: Remove all aliases (use with care!)	unalias command
\$ uname	Display technical information about current system	-a: Display all basic information	uname -a
\$ unzip	Unzip a Windows-compatible Zip file	-l: Display archive content but don't actually unzip	unzip myfile.zip

Command	Description	Typical Command Options	Examples of Use
\$ uptime	Display uptime for system, as well as CPU load average and logged-in users		uptime
\$ vi *	Text editor program		vi
\$ wc	Count the number of words in a file		wc <i>myfile.txt</i>
\$ whatis	Display one-line summary of specified command		whatis <i>command</i>
\$ whereis	Display information on where a binary command is located, along with its source code and man page (if applicable)	-b: Return information only about binary programs	whereis -b <i>command</i>
\$ xhost	Configure which users/systems can run programs on the X server	+: When followed by a username and/or system name, gives the user/system permission to run programs on the X server; when used on its own, lets <i>any</i> user/system use the X server -: Opposite of +	xhost +
\$ xinit	Start elementary GUI session (when not already running a GUI)		xinit
\$ zip	Create Windows-compatible compressed Zip files	-r: Recursive; includes all subdirectories and files therein -u: Updates Zip with specified file -P: Encrypts Zip with specified password -v: Verbose; display more information -#: Set compression level (from 0, which is no compression, to 9, which is highest)	zip -r <i>myfile.zip mydirectory</i>
\$ zipgrep	Searches inside Zip files for specified text string		zipgrep " <i>search phrase</i> " <i>myfile.zip</i>



Getting Further Help

So you've read through this book and have a good working knowledge of Linux. Ubuntu is running exactly as you want it to, and things are going okay. But then you hit a brick wall. Perhaps you want to perform a task but simply don't know how. Or maybe you know roughly what you need to do but don't know the specifics. Although, this book tries to be as comprehensive as possible, it can't cover every eventuality.

You need to find some help, but where do you turn? Fortunately, many sources of information are available to those who are willing to help themselves. Linux contains its own series of help files in the form of man and info pages, and these are good places to start. In addition, some programs come with their own documentation. If neither of these sources provides the help you need, you can head online and take advantage of the massive Linux community around the world.

Read the Manual!

Before asking for help online, it's important that you first attempt to solve your problems by using Linux's built-in documentation. If you go online and ask a question so simple that it can be answered with a little elementary research, you might find people reply with "RTFM." This stands for Read the Freaking Manual. In other words, do some basic research, and then come back if you're still stuck.

It's not that people online don't want to help. It's that they don't like people who are too lazy to help themselves and expect others to do the hard work for them. Although not all Linux people you encounter will take such a hard line, doing a little homework first can provide answers to a lot of questions, removing the need to ask others. This is particularly true when it comes to the fundamentals.

Documentation typically comes in three formats: man pages, info pages, and README files.

Man Pages

Man pages are the oldest form of Unix documentation. In the old days, once an individual had created a piece of software, he would write a brief but concise man page in order to give others a clue as to how to operate it. The programmer would come up with a few screens of documentation that could be called up from the command prompt. This documentation would outline what the software did and list all the ways in which it could be used.

Nowadays, depending on the software package, man pages are usually created by technical writers, but the concept of providing essential information still applies. Man pages under Linux provide all the information you need about how to use a particular command or piece of software.

Sounds great, doesn't it? Alas, there's a problem: man pages are written by software engineers *for* software engineers. They expect you to already understand the technology being discussed. This is illustrated very well by the man page for `cdrecord`, software that can be used to burn CD images to disc. You can view this man page by typing `man cdrecord` at the command prompt.

The first line of the man page states, "Cdrecord is used to record data or audio Compact Discs on an Orange Book CD-Recorder or to write DVD media on a DVD-Recorder."

Most of that is clear, but what do they mean by "Orange Book"? They don't explain. (If you're curious, head over to http://searchstorage.techtarget.com/sDefinition/0,,sid5_gci503648,00.html.)

Further down in the man page, you see, "Cdrecord is completely based on SCSI commands ... Even ATAPI drives are just SCSI drives that inherently use the ATA packet interface as [a] SCSI command transport layer."

What's SCSI, or ATAPI for that matter? Again, the man page doesn't explain. (They're methods of interfacing with storage devices attached to your computer.)

But why should man pages explain as they go along? Their function is to describe how to use a piece of software, not to provide a beginner's introduction to technology. If they did that, a single man page could run to hundreds of pages.

In other words, man pages are not for complete beginners. This isn't always the case and, because Linux sees widespread usage nowadays, man pages are sometimes created with less knowledgeable users in mind. But even so, the format is inherently limited: man pages provide concise guides to using software. Luckily, there are some tips you can bear in mind to get the most from a man page. But before you can use those tips, you need to know how to read a man page.

How to Read a Man Page

To read a man page, you simply precede the command name with `man`. For example, to read the man page of `cdrecord`, a piece of software used to write ISO images to CD-R/RW discs, type the following command:

```
man cdrecord
```

This opens a simple text viewer with the man page displayed. You can move up and down line by line with the cursor keys, or move page by page using the Page Up and Page Down keys (these are sometimes labeled Pg Up and Pg Down). You can search by hitting the forward slash key (`/`). This will highlight all instances of the word you type. You can search for other examples of the word in the document by hitting the `n` key.

The average man page will include many headings, but the following are the most common:

Name: This is the name of the command. There will also be a one-sentence summary of the command.

Synopsis: This lists the command along with its various command options (sometimes known as *arguments* or *flags*). Effectively, it shows how the command can be used. It looks complicated, but the rules are simple. First is the command itself. This is in bold, which indicates it is mandatory. This rule applies to anything else in bold: it must be included

when the command is used. Anything contained within square brackets ([]) is optional, and this is usually where you will find the command options listed. A pipe symbol (|) separates any command options that are exclusive, which means that only one of them can be used. For example, if you see [apple|orange|pear], only one of apple, orange, or pear can be specified. Usually at the end of the Synopsis listing will be the main argument, typically the file(s) that the command is to work on and/or generate.

Description: This is a concise overview of the command's purpose.

Options: This explains what the various command options do, as first listed in the Synopsis section. Bearing in mind that command options tell the software how to work, this is often the most useful part of the man page.

Files: This lists any additional files that the command might require or use, such as configuration files.

Notes: If this section is present (and often it isn't), it sometimes attempts to further illuminate aspects of the command or the technology the command is designed to control. Unfortunately, Notes sections can be just as arcane as the rest of the man page.

See Also: This refers to the man pages of other commands that are linked to the command in question. If a number appears in brackets, this means the reference is to a specific section within the man page. To access this section, type: `man <section no> command`.

Although there are guidelines for the headings that should appear in man pages, as well as their formatting, the fact is that you may encounter other headings, or you may find nearly all of them omitted. Some man pages are the result of hours if not days of effort; others are written in ten minutes. Their quality can vary tremendously.

Tips for Working with Man Pages

The trick to quickly understanding a man page is decoding the Synopsis section. If you find it helps, split the nonobligatory command options from the mandatory parts. For example, `cdrecord`'s man page says that you *must* specify the `dev=` option (it's in bold), so at the very least, the command is going to require this:

```
cdrecord dev=X filename
```

Then you should skip to the Options section and work out which options are relevant to your requirements. While you're there, you'll also need to figure out what the `dev=` command option requires.

Although the command options contained in square brackets in the Synopsis section are, in theory, nonobligatory, the command might not work satisfactorily without some of them. For example, with `cdrecord`, I use the `-speed` command option, which sets the burn speed, and also the `-v` option, which provides verbose output (otherwise, the command runs silently and won't display any information on screen, including error messages!).

Another handy tip in decoding man pages is understanding what standard input and standard output are. In very simple terms, standard input (stdin) is the method by which a command gets input—the keyboard on most Linux setups. Standard output (stdout) is where the output of a command is sent, which is the screen on most Linux setups. (See Chapter 15 for more details about standard input and standard output.)

Often, a man page will state that the output of a command will be sent to standard output. In other words, unless you specify otherwise, its output will appear on screen. Therefore, it's necessary to specify a file to which the data will be sent, either by redirecting the output (see Chapter 17), or by specifying a file using a command option. For example, the `mkisofs` command can be used to create ISO images from a collection of files for subsequent burning to CD. But unless the `-o` option is used to specify a filename, `mkisofs`'s output will simply be sent to standard output—it will appear on the screen.

Finally, here's the best tip of all for using man pages: don't forget that `man` has its own man page. Simply type `man man`.

Info Pages

Man pages date from the days of relatively primitive computers. Back then, most computers could only display page after page of text, and allow the user to scroll through it. In addition, memory and disk space were scarce, which is why some man pages are incredibly concise—fewer words take up less memory!

The *Texinfo* system is a valiant by the GNU Project attempt to overcome the shortfalls of man pages. Often, this is referred to as *info*, because that's the command used to summon Texinfo pages (normally, you type `info command`).

For starters, info pages are more verbose than the equivalent man pages, and that gives the author more space to explain the command or software. This doesn't necessarily mean that info pages are easier to understand, but there's a better chance of that being the case.

Secondly, info pages contain hyperlinks, just like web pages. If you move the cursor over a hyperlinked word, which is usually indicated by an asterisk (*), you can proceed to a related page. In a similar sense, pages are linked together so that you can move back and forth from topic to topic.

The bad news is that the man page system is far more popular and established than Texinfo. If a programmer creates a new application, for example, it's unlikely he'll bother with an info page, but he will almost certainly produce a man page.

In fact, in many cases, typing `info command` will simply bring up the man page, except in the software used to browse info pages.

However, nearly all the GNU tools are documented using info pages, either in their own pages or as part of the `coreutils` pages. For example, to read about the `cp` command and how to use it, you can type this:

```
info coreutils cp
```

To browse through all sections of the `coreutils` pages, type this:

```
info coreutils
```

Because man pages are so established, everyone expects to find one for every utility. So most utilities that have info pages will also have man pages. But in such a case, the man page will state near the end that the main documentation for the utility is contained in an info page and you may find it more fruitful to use that instead.

Navigating through info pages is achieved via the keyboard and is something of an art. But, as you might expect, there's a user-friendly guide to using info: just type `info info`. Remember that words preceded with an asterisk are hyperlinks, and you can jump from link to link using the Tab key.

README Files and Other Documentation

Some programs come with their own documentation. This is designed to give users the information they need to get started with the program (as opposed to the man page, which is a concise and complete guide to the software). Alternatively, program documentation sometimes gives a brief outline of the program's features.

The files are usually simple text, so they can be read in any text editor or word processor, and are normally called README. Under Ubuntu, these documents are usually stored in a program-specific directory within `/usr/share/doc` (although a small minority of programs use `/usr/doc`).

Not all programs are friendly enough to provide such documentation, but even so, you'll still find a directory for the software in `/usr/share/doc`. This is because the software might also come with a getting started guide written by the Ubuntu package maintainer. Such guides detail specifics of using the software under Ubuntu, such as where configuration files are located or how the program interoperates with other software on the system. Sometimes, this documentation is written by a Debian package maintainer because nearly all Ubuntu software has its origins in the Debian project (www.debian.org).

In addition, the directory will probably contain copyright information, explaining the software license used by the software, as well as a CHANGelog, which is a text file listing features that have been added to each release of the software. The directory might contain some other files, too, detailing where to send information about bugs, for example.

Viewing the README documentation is easy. For example, for the Gedit text editor, you could type this:

```
cd /usr/share/doc/gedit
less README
```

Sometimes, the README documentation is in a compressed tarball, in which case it will have either a `.tar.gz` or a `.tar.bz2` file extension. However, `less` is clever enough to realize this and extract the document for reading.

Getting Help Online

If you can't figure out the answer by referring to the documentation, then there's little choice other than to look online. Linux benefits from a massive community of users, all of whom are usually willing to help each other.

Usually, the best way of getting help is to visit a forum. Here, you can post messages for others to reply to. Alternatively, you might choose to sign up for a mailing list. This is a way of sending e-mail to several hundreds, if not thousands, of people at once. Any individual can then reply. Mailing lists often have the benefit of allowing personal attention and interaction, but this comes at the expense of each subscriber receiving a whole lot of mail.

Forums

The official Ubuntu project forums are located at www.ubuntuforums.org. You'll find forums for just about every need, from security to beginner's issues, but by far, the most popular is the one devoted to the current release of Ubuntu. At the time of writing, this was Ubuntu 5.10 Support (GNOME).

Before you can post, you need to register by providing an e-mail address. This is designed to keep down the quantity of unwanted junk postings to the forum.

You might think it fine to post a new question immediately after registering, but don't forget the simple rules mentioned at the beginning of this appendix: if you don't do elementary research first and try to solve your own problem, you may elicit a hostile response from the other posters, especially if your question is one that comes up time and time again, and has been answered several times.

So, first make use of the comprehensive search facility provided with the forums. For example, if you're looking for advice on getting a Foomatic D1000 scanner working, use this as a search term and see what comes up. The chances are that you won't be the first person who has run into problems with that piece of hardware, and someone else may have already posted a solution.

Often, you'll need to read the full thread to find an answer. Someone may start by asking the same question as you but, with the help and guidance of the forum members, they might find a solution, which they then post several messages later.

In addition, some individuals write their own HOWTO guides when they figure out how to do something. These are normally contained in the Customization Tips & Tricks forum.

If you're unable to find a solution by searching, then consider posting your own question. Keep your question simple, clear, and concise, because no one likes reading through acres of text. If possible, provide as many details about your system as you can. You will almost certainly want to provide the version number of the Linux kernel you're using, for example. You can find this version number by typing the following in a GNOME Terminal window:

```
uname -sr
```

In addition, any other details you can provide may prove handy. If you're asking about hardware, give its entire model name and/or number. Don't just ask for help with a Foomatic scanner. Ask for help with a Foomatic D1000 scanner, model number ADK1033, Revision 2. If you're asking about a piece of software, provide its version number (click Help ► About).

Sometimes in their replies, other forum members may ask you to post further details or to provide log files. If you don't understand the question, simply ask the poster to give you more details and, if necessary, instructions on what to do. Just be polite. Explain that you're a newbie. If you think the question is extremely obvious, then say so—apologize for asking what may be a stupid question, but explain that you've tried hard to answer it yourself but have failed. Don't forget that the Ubuntu forums include the Absolute Beginner Talk forum, where fundamental questions are asked all the time.

Mailing Lists

Using the forum's search function also has the advantage of searching the archives of the mailing lists.

Mailing lists have a number of advantages and disadvantages. The advantages are that a mailing list provides an excellent way to learn about Ubuntu. All you have to do is read through the e-mail messages you receive in order to partake of a constant information drip-feed. In addition, some mailing lists are designed to make public announcements, so you'll find it easy to learn about the latest happenings in the Ubuntu community.

Mailing lists also have a terrific sense of community. They offer a neat way of getting to know other Ubuntu users and talking to them. E-mails often drift off topic into humor and general discussion.

The disadvantages of mailing lists are that you can easily receive in excess of 200 messages a day, depending on which mailing list you join. Even if you have a moderately fast Internet connection, that quantity of messages can take a long time to download. In addition, you'll need to sort out any personal or business e-mail from the enormous quantity of mailing list traffic (although the mailing list messages usually have the list title in square brackets in the subject field; you can therefore create a mail rule that sorts the mail according to this).

You can sign up to the Ubuntu mailing lists at www.ubuntulinux.org/community/lists/document_view.

Other Official Sites

The Ubuntu Document Storage Facility (http://doc.gwos.org/index.php/Main_Page) is a guide put together by the individuals who maintain the Ubuntu forums. It contains a lot of information culled from the forums, not least of which is the Hardware Compatibility Guide. This lists hardware that is known to work with Ubuntu, or at least hardware that can be made to work with Ubuntu with a little effort, which often is also detailed.

The official Ubuntu web site, www.ubuntu.org, contains a Documentation section that features FAQs and a glossary of terms. Simply click the Documentation link on the right side of the main web page.

In addition, you might want to take a look at the Ubuntu wiki: <https://wiki.ubuntu.com>. This contains a whole world of fascinating information about Ubuntu, but can be somewhat difficult to navigate. Try starting with the Documentation pages: <https://wiki.ubuntu.com/UserDocumentation>.

Third-Party Sites

Of course, the Ubuntu project doesn't have a monopoly on sites that discuss Ubuntu. Several third-party web sites are worth at least an occasional visit, and other forum web sites are devoted to Linux.

One I visit on a regular basis is the Ubuntu blog: <http://ubuntu.wordpress.com>. This is written by a relative newcomer to Ubuntu who has chosen the Ubuntu distribution for his computer. It details discoveries that Carthik Sharmahas made about Ubuntu, such as configuration options, online sources of software, how to overcome hardware issues, and so on.

In addition, I like to visit <http://linuxhelp.blogspot.com>, which is a similar blog written by a Linux user who uses Ubuntu and likes to share tips and techniques.

Perhaps the king of third-party Ubuntu sites is Ubuntu Guide: <http://ubuntuguide.org>. This contains brief instructions on how to do a variety of common tasks under Ubuntu, such as installing certain types of software or administering particular hardware. It covers a lot of the same ground as this book, but is still worth investigating if you wish to browse through some excellent tips and advice.

Finally, one of the best Linux forums and general advice sites can be found at www.linuxquestions.org. This has a forum dedicated specifically to Ubuntu, but also contains hundreds more devoted to just about every aspect of Linux, including forums for beginners.



Other Versions of Ubuntu

The nature of open-source software encourages diversification, so it should come as no surprise that, along with the main release of Ubuntu, there are several offshoot projects. All of these take the standard Ubuntu release as a base but tweak it to various degrees, perhaps to use a different graphical desktop system or to cater to different types of users.

This appendix lists a number of such projects. Perhaps it goes without saying that all these projects are free of charge, just like Ubuntu, and are only a download away.

Kubuntu

The standard Ubuntu release, as supplied with this book, relies on the GNOME Desktop Project for its graphical interface (see www.gnome.org). Many other desktop projects exist in the wider Linux world, but perhaps the only one that ranks alongside GNOME in terms of popularity is the K Desktop Environment project (www.kde.org), usually referred to as KDE. Kubuntu is simply a version of Ubuntu that eschews GNOME in favor of KDE.

Kubuntu (www.kubuntu.org) retains the same philosophy as Ubuntu, in both its humanitarian aims of being available to all as well as its more pragmatic aspects, such as always including the latest versions of applications. It also shares many technical features, such as the use of `sudo` to invoke superuser powers.

The main difference is the software bundled with this variant. When it comes to e-mail, KDE's KMail program is used instead of Evolution, for example, and Konqueror is used for browsing instead of Firefox (although, of course, Evolution and Firefox can easily be installed via the Synaptic Package Manager after Kubuntu has been installed). The system configuration software is radically different, too, with several KDE tools used instead of the GNOME software described in this book.

To try Kubuntu, you can either download an ISO image for burning to CD-R/RW from <http://releases.ubuntu.com>, or simply search for and install the `kubuntu-desktop` package using the Synaptic Package Manager. This will then install all the Kubuntu components. Once it's installed, you can opt to boot into Kubuntu by click the Session button on the login screen and selecting the KDE entry from the list.

Edubuntu

The Ubuntu philosophy is to make an operating system accessible by everyone, no matter who they are or where they live in the world. Of course, young people are included in this vision, and Edubuntu (www.edubuntu.org) is a version of Ubuntu specifically geared towards their needs. It's a rendition of the standard Ubuntu release that's been bolstered by many educational software titles. In addition, it has a friendlier and simpler interface that's designed to appeal to youngsters.

The chief education titles are KDE Edutainment Suite, GCompris, and Tux4Kids. KDE Edutainment Suite (<http://edu.kde.org>) contains games involving mathematics, languages, science, and other miscellaneous topics. Teachers might also appreciate the inclusion of KEduca, a program designed to create form-based tests and exams.

GCompris (<http://gcompris.net>) is for children of kindergarten age and introduces them to computer use as well as elementary math and reading skills. Tux4Kids includes a typing tutor, publishing program, and a handful of games geared towards math and literacy skills.

The main benefit of Edubuntu for educational establishments is that it's both free of charge and comes with the same kind of update support as Ubuntu, meaning security updates will be available for at least 18 months after the initial release.

Ubuntu Server

The standard version of Ubuntu, as supplied with this book, is designed primarily for desktop computers (sometimes referred to as *workstations*). It can be used as an elementary server operating system, but for more demanding tasks, additional configuration is necessary. A much better choice for practically any kind of server need is Ubuntu Server (<http://releases.ubuntu.com/ubuntu-server>).

Servers are computers whose sole purpose in life is to both store and make available data or resources for other computers via a network connection. For example, a *file server* is a central repository of files that's often used in offices or other establishments. A *web server* is a computer that's connected to the Internet and provides web pages and other data for web surfers. You might also have heard of e-mail servers, FTP servers, and others.

Many Linux server operating systems can be difficult to use and hard to set up. Ubuntu Server overcomes these issues while attaining hardware certification and maintaining quality-assurance levels. Certification is an award given out by manufacturers that guarantees their hardware will work under certain operating systems. The Ubuntu Server project's goal is to achieve this for most common types of server hardware. Quality assurance is a way of guaranteeing that the operating system can be relied on to be stable.

In terms of ease of use, the Ubuntu Server project aims to create a server operating system that can be used by all levels of users, not just experts. For those who have little experience of Linux but need to create a server system, it's an excellent choice.

UBUNTU UPDATES

The Ubuntu Foundation is unique in the world of Linux distributions in that it works to a six-month release cycle. This means a new version of Ubuntu comes out every half year. Most other Linux distros are released yearly at best, although it's not unusual for several years to pass before the release of a new version of some distros.

The version of Ubuntu offered with this book is 5.10, code-named Breezy Badger (each release receives a humorous code name, usually related to animals). This was the latest release available as this book was being written. By the time you read this, it's very likely there will be a new release of Ubuntu available.

You can update to the latest version of Ubuntu in two ways. The first is to download the ISO image of the latest release from <http://releases.ubuntu.com> and burn it to CD. This can be done within Ubuntu by right-clicking the ISO file and selecting Write to Disc (ensure you select a relatively slow burn speed in the dialog box that appears; setting a fast write speed will probably result in a unusable disc). Then, when the CD has been burnt, simply reinsert it into the CD/DVD-ROM drive. When you insert the CD, Ubuntu will ask if you want to upgrade to the version of the CD.

The other way of upgrading is to open a GNOME Terminal window (Applications ► Accessories ► Terminal) and type the following:

```
sudo apt-get dist-upgrade
```

This will then download all the packages for the latest release of Ubuntu and attempt to update your system.

Xubuntu

While GNOME and KDE dominate the desktop interface landscape of Linux, and they are used in the main Ubuntu and Kubuntu releases, respectively, other projects take a different approach to the graphical desktop. XFCE (www.xfce.org) is one of these. It's a streamlined desktop that retains the good looks of GNOME but is much smaller in terms of memory footprint. This means that a system running XFCE is faster than an equivalent GNOME system. It also means that XFCE can be used on many older computers that don't have the powerful hardware we take for granted nowadays and which would struggle with the latest GNOME and KDE releases of Ubuntu.

Because a key component of the Ubuntu Foundation's philosophy is to create an operating system that can be used by anyone, regardless of where they are in the world, a version of Ubuntu that can run on older hardware makes a lot of sense. It's unlikely that less developed countries will have access to the latest expensive computer hardware, for example.

Xubuntu (<https://wiki.ubuntu.com/Xubuntu>) is simply a version of Ubuntu that replaces GNOME with XFCE. Despite XFCE's claim to be lightweight, it still offers a high degree of usability and shouldn't be seen as a second-best choice for stripped-down hardware. It's certainly worth trying out. It also uses many modern GUI aspects we take for granted, such as theming (see Chapter 10) and font antialiasing.

As with Kubuntu, you can upgrade to Xubuntu by either burning a CD or by using APT via the following command:

```
sudo apt-get install xubuntu-desktop
```

Alternatively, just search for `xubuntu-desktop` using the Synaptic Package Manager. Once it's installed, select XFCE after clicking the Session button on the login screen.

Other Ubuntu Projects

If you have an Apple Macintosh computer that uses a G3, G4, or G5 processor, you can download the PowerPC version of Ubuntu. This is identical to the standard Ubuntu release in practically every way. The release works on iBook and PowerBooks, too.

There's also a version of Ubuntu for 64-bit processors. This includes those that use the AMD64 instruction set, such as AMD's Athlon 64 chips and Intel Pentium 4 processors that feature EMT64 support (Intel's earlier Itanium range of 64-bit chips aren't supported). Once again, this release is practically identical to the main Ubuntu release. The only difference is that the software has been compiled for 64-bit processors.

Ubuntu offers a number of "live" distros. These are versions of Ubuntu that run from the CD and don't install to the hard disk. Live distros are handy for introducing Ubuntu to friends who aren't yet ready to abandon Windows. A disadvantage is that Ubuntu runs a little slowly because the data must be read from the CD/DVD-ROM drive, rather than from the hard disk. You can find live versions of practically all the main Ubuntu releases, including Ubuntu, Kubuntu, and PowerPC.

nUbuntu (<http://nubuntu.org>) is a version of Ubuntu geared towards those who work in the computer security field and includes many network security testing tools. It can also be used by power users. It features the FluxBox desktop interface (<http://fluxbox.sourceforge.net>) instead of GNOME or KDE.

All of these releases, with the exception of nUbuntu, are available for download as ISO images at the Ubuntu download site: <http://releases.ubuntu.com>. To burn any ISO to CD, insert a blank CD, right-click the ISO file, and select Write to Disc. Make sure you select a slow burn speed in order to obtain a usable disc.

Index

■ Symbols

- (three dashes), 226
- \$ (dollar symbol), 207
- \$ command, vi command mode, 246
- \$PATH variable, 267
- %CPU column, top program, 259
- %MEM column, top program, 259
- & (ampersand symbol), 264
- * (asterisk), 532
- * wildcard, 251
- | (pipe symbol), 254
- | symbol, 273
- ~ (tilde), 474
- < (left angle bracket), 254
- = (equal sign), Calc, 355
- > (angle bracket), 274
- > (right angle bracket), 254
- ! (exclamation mark), 271, 484
- / (forward slash), 183
- . (period), 235, 463
- ? (question mark), 242
- / (slash), 242
- ! command, sftp, 485
- / command, vi command mode, 246
- ? command, vi command mode, 246

■ A

- a command option, 254, 463
- A command, vi command mode, 246
- AbiWord word processor, 428
- acceleration, 146
- Acceleration setting, 146
- Access Path field, 193

- Accessibility option, OpenOffice.org, 335
- Accessibility tab, GNOME Display Manager, 149
- Accessibility tools, 150
- accessing computers remotely. *See* remote access to computers
- Account tab, 438
- acpid service, 448
- acpi-supply, 448
- Active Directory/Global Address List, 103
- adaptername, 89
- Add Directory button, 466
- Add File button, 466
- Add New button, 475
- Add to Panel option, 156
- Add Wallpaper button, 143
- Address Cards view, Contacts mode (Evolution), 394
- adm system group, 440
- admin system group, 441
- administrator account, 65
- Adobe Photoshop 6, 409
- Advanced option, 58
- Advanced Packaging Tool (APT), 417, 459
- Advanced tab, 438
- age of Linux, 5–6
- Airbrush tool, GIMP, 305
- Alcatel SpeedTouch series USB DSL modem, 92, 95–98
- alias command, 515
- alsa service, 448
- alsamixer command, 515
- alsa-utils service, 448
- Altering Permissions section, 226

- AMD64 instruction set, 540
 - ampersand symbol (&), 264
 - anacron, 449, 476–477
 - angle bracket (>), 274
 - animating slides, 366–367
 - Animation filters, GIMP, 315
 - Ant screensavers, 118
 - antialiasing, 144
 - antivirus software, 6, 129
 - API (Application Programming Interface), 401
 - apmd service, 449
 - Appearance option, OpenOffice.org, 335
 - applets, 156
 - Application category, OpenOffice.org, 332
 - Application Programming Interface (API), 401
 - applications. *See* programs
 - Applications menu, 67, 71, 155, 161
 - Apply button, Synaptic Package Manager toolbar, 422
 - apropos command, 515
 - APT (Advanced Packaging Tool), 417, 459
 - APT and package repositories
 - categories of repositories, 418–419
 - overview, 417–418
 - repository sections, 419–420
 - APT tools
 - installing and removing packages, 427–428
 - overview, 426–427
 - querying packages and repositories, 429–430
 - updating system, 428
 - apt-cache command, 516
 - apt-cache depends packagename |less command, 429
 - apt-cache search packagename command, 429
 - apt-cache show packagename |less command, 429
 - apt-get command, 427, 431, 516
 - Archive Manager, 177
 - Artistic filters, GIMP, 314
 - As Character option, Writer, 346
 - Assistive Technology Support, 150
 - asterisk (*), 532
 - at command, 477–478
 - atd service, 449
 - ATI 3D cards, 116
 - ATI driver, installing, 117
 - atq job number, 478
 - Audio file types, 191
 - audio group, 440
 - audio playback programs, 171–172
 - audio system group, 440
 - AutoCaption option, OpenOffice.org Writer, 337
 - autocomplete function, 268
 - autocompletion, BASH shell
 - with files and paths, 268
 - overview, 267
 - viewing available options, 268–269
 - Autohinting, 145
 - automatic partitioning tool, 39
 - Available Backups drop-down list, 468
 - Away icon, Gaim, 178
- B**
- b command, vi command mode, 246
 - ba commands, 268
 - background (bottom) color box, 305
 - Background option, OpenOffice.org Writer/Web, 337
 - backing up data
 - choosing data to back up, 463–464
 - from command line
 - compressing tar archives, 469–470
 - creating tar files, 469
 - extracting files from tar archive, 470
 - overview, 468
 - saving file to CD-R/RW, 471
 - viewing tar archive information, 470

- e-mail, 30–31
- making notes, 31–32
- overview, 29–30, 463
- using SBackup
 - changing backup file destination, 467
 - changing time period between backups, 467–468
 - excluding files and folders, 466
 - including files and folders, 466
 - overview, 464–466
 - restoring data via SBackup, 468
- Backports repository, 419
- backslash (\), 183
- Backspace key, 249
- Backup Properties dialog box, 465
- bang, 271
- Base (database program), 167
 - creating database
 - adding controls to form manually, 378–379
 - adding table, 375–377
 - creating form, 377–378
 - editing database table, 379–380
 - overview, 375
 - getting started with, 373–375
 - overview, 373
 - using database, 380–381
- Base program window, 374
- BASH shell, 204
 - autocompletion
 - with files and paths, 268
 - overview, 267
 - viewing available options, 268–269
 - command history, 271–272
 - commands, 257, 515–527
 - getting help, 209–210
 - keyboard shortcuts
 - overview, 269
 - for system control, 270–271
 - for working in BASH, 269–270
 - overview, 203–208, 267
 - pipng and directing output
 - of commands, 272–274
 - overview, 272
 - redirecting output, 274–275
 - reasons for using, 205–206
 - running programs, 208–209
 - running via virtual console, 210–211
 - when to use, 206–207
 - working with files
 - changing and creating directories, 217
 - copying files, 214–215
 - deleting files, 216–217
 - listing files, 213–214
 - moving files, 215–216
 - overview, 211–213
- BASIC category, OpenOffice.org, 332
- Basic Fonts (Western) option, Writer, 336
- BBC (British Broadcasting Corporation), 295
- benchmarking hard disk, 452–453
- benefits of Linux, 8–9
- bin directory, 230
- binary file, 414
- binary value, 167
- BIOS, 34, 47–48, 443
- BitTorrent file sharing application, 133
- Bksp key, GNOME Calculator, 175
- blank screens, 51
- Bluetooth, 447
 - overview, 118–119
 - pairing devices, 119
 - transferring files between devices, 120–121
 - using keyboard or mouse, 122–123
- Bluetooth File Sharing, 120
- Bluetooth Manager, 120
- bluez-uti, 449
- Blur filters, GIMP, 311
- Blur/sharpen tool, GIMP, 305

Boolean value, 167
 Boot Device Priority, 34
 boot directory, 230
 boot loaders, 26, 63–64
 boot sectors, 26
 Boot Sequence, 34
 booting

- for first time, 45–46
 - cutting and pasting text, 74
 - exploring desktop, 66–71
 - logging in, 64–65
 - overview, 63
 - running programs, 71–72
 - starting up, 63–64
 - using mouse, 73–74
 - working with virtual desktops, 72–73
- speeding up
 - initialization and run-level services, 446–452
 - overview, 443–445
 - reducing boot menu delay, 445

/boot/grub/menu.lst file, 445
 bootlogd service, 449
 Borders icon, Writer, 349
 Breezy Badger, 420, 539
 Brightness/Contrast menu option, 308
 British Broadcasting Corporation (BBC), 295
 browser dialog box, GIMP, 302
 Brush tool, GIMP, 305
 BT Voyager USB ADSL modem, 92–95
 buffers, GIMP, 306
 bugs, and proprietary drivers, 78
 build-essential metapackage, 431
 Burn and dodge tool, GIMP, 305
 Buttons tab, Mouse Preferences dialog box, 146
 By Company view, Contacts mode, Evolution, 394
 Bytecode Interpreting, 145
 bzip2 program, 177, 469, 517
 bzip2recover command, 517

C

c command option, 469
 c hdparm command option, 454
 cable broadband, 81
 cache of package files, emptying, 459–460
 cal command, 517
 Calc (spreadsheet program), 165–166

- creating charts, 359–361
- deleting and inserting data and cells, 356
- entering and formatting data, 355–356
- overview, 355
- sorting data, 359
- using filters, 361
- working with formulas
 - overview, 356–357
 - summing figures, 358
 - using Function Wizard, 357

Calculate option, Calc, 337
 Calculator, 162, 175–176
 Calendar view, Evolution, 168
 Calendars mode, Evolution, 385

- adding or editing diary entry, 397–398
- overview, 396
- specifying appointment types, 396

camera. *See* digital cameras
 case sensitivity, 184
 cat command, 232, 241, 274, 517
 Category drop-down list, Calc, 357
 cd command, 272, 484, 517
 cdparanoia * command, 517
 CD player, 162
 CD/DVD burning, 50, 162, 173
 cdrecord command, 275, 471, 517
 /cdrom directory, 230, 232
 cdrom system group, 440
 CD-R/RW, saving backup file to, 471
 CDs (compact discs)

- audio, listening to, 283–284
- creating own, 288–289

- ripping music from
 - adding MP3 support to Sound Juicer, 286–287
 - choosing format, 285–286
 - overview, 285
 - process, 287–288
- cells, in Calc, 356
- centered option, Style drop-down list, 144
- central file server, 220
- cdisk command, 460, 517
- Change Icon, 333
- Changes option
 - Calc, 337
 - Writer, 337
- Character option, Writer, 344
- charts, Calc, 359–361
- chgrp command, 517
- child processes, 261
- chmod +w README command, 249
- chmod command, 227, 518
- chown command, 468, 518
- chroot command, 518
- chvt command, 518
- Cinelerra software suite, 299
- ClamAV, 129
- clear command, 518
- Clearlooks theme, 140
- clock, 68
- cloning, in GIMP, 309–310
- coder-decoder, 280
- coding/decoding, 90
- color, correcting, in GIMP, 308–309
- Color filters, GIMP, 312
- color ink-jet models, 107
- Color picker tool, GIMP, 304
- Color region selection tool, GIMP, 303
- color-coding, within GNOME Terminal, 221
- Colors option, OpenOffice.org, 335
- Combine filters, GIMP, 315
- COMMAND column, top program, 259
- !command command, sftp, 485
- command history, BASH shell, 271–272
- command line
 - adding and deleting users at, 439–441
 - backing up data from
 - compressing tar archives, 469–470
 - creating tar files, 469
 - extracting files from tar archive, 470
 - overview, 468
 - saving file to CD-R/RW, 471
 - viewing tar archive information, 470
- Command mode, vi text editor, 245–246
- command prompts, 207
- command-line mode, 249
- Command-Line mode, vi text editor, 247–248
- command-line prompts, 203
- command-line shell, 204–206, 229, 243, 250
- commands, piping and directing output of, 272–274
- comma-separated values (CSV), 166, 321, 349
- community, 9–10
- compact discs. *See* CDs
- compatibility, 8–9, 321–323
- Compatibility option, Writer, 337
- compiling a kernel, 458
- composing e-mail messages, 388
- Compressed file types, 191
- compressing tar archives, 469–470
- computer memory cards, 111
- Computer window, 197
- configuration dialog box, GIMP, 174
- Configuration Editor program, 155
- configuration options, 334
- configuration software, 20
- configuration tools, 21
- configure script, 433

- configuring
 - Firestarter
 - overview, 132–133
 - setting inbound rules, 133–134
 - setting outbound rules, 134–135
 - turning off diagnostic services, 135–136
 - monitor, 44–45
 - Wine, 403–406
 - Connections option, Base, 338
 - Contact List Editor pane, Evolution, 396
 - Contacts manager/calendar, 162
 - Contacts mode, Evolution, 384
 - adding or editing contact information, 395
 - creating contact list, 396
 - overview, 394
 - Contacts view, Evolution, 168
 - Content Scrambling System (CSS), 296
 - context menu, Writer, 344
 - Contiguous regions selection tool, GIMP, 303
 - Control Panel function, Windows, 70
 - controls, 141
 - Controls category, OpenOffice.org, 332
 - copying
 - files, BASH shell, 214–215
 - Windows fonts, 323–325
 - copyleft, 14
 - core fonts, 323
 - costs of Linux, 8
 - cp command, 214, 518
 - cp program, 237
 - cpio command, 468
 - CPU Frequency Scaling Monitor, GNOME, 102
 - crashes, 8
 - Create Form, Base, 377
 - Create Launcher dialog box, 152
 - cron, 449, 473
 - crontab, 473–474, 477, 518
 - creating scheduled task with, 474–475
 - editing with GUI, 475
 - overview, 473
 - Crop tool, GIMP, 304, 309
 - cropping, in GIMP, 309–310
 - CrossOver Office, 299
 - CSS (Content Scrambling System), 296
 - CSV (comma-separated values), 166, 321, 349
 - CUPS Printer (IPP), 108
 - cupsys service, 449
 - Currency icon, Calc, 355
 - cursor keys, 33, 36, 64, 245
 - Cursors tab, Mouse Preferences dialog box, 146
 - Custom Animation dialog box, 366–367
 - Custom backup mode, 465
 - Custom button, Synaptic Package Manager toolbar, 421
 - Customize dialog box, 333
 - cutting and pasting text, 74
- D**
- D command, vi command mode, 245
 - Data category, OpenOffice.org, 332
 - Database pane, Base, 374, 380
 - Database Wizard, 373
 - databases. *See* Base
 - Databases option, Base, 338
 - date command, 518
 - db command, vi command mode, 245
 - dbus service, 449
 - dd command, vi command mode, 245
 - deactivating
 - prelinking, 457–458
 - spell-checking, 346
 - .deb package, 424
 - Debian distribution, 22
 - Debian package maintainer, 533

- Debian package-management system, 415, 430
- DeCSS software, 296
- Default Colors option, Chart, 338
- default mode, GNOME Calculator, 175
- Delete key, Calc, 356
- Delete This Panel option, 156
- deleting
 - data and cells, in Calc, 356
 - e-mail messages, 386–387
 - files, BASH shell, 216–217
 - menus, 155
 - Windows partition, 39
- denial-of-service (DoS) attacks, 135
- dependencies, 416, 425, 431–432
 - management, 416–417, 423
 - resolution, 417
- Description heading, man page, 531
- desktop, 66
 - adding shortcuts
 - creating launcher, 152–153
 - creating link, 153–154
 - overview, 152
 - applets, 156
 - elements, 68–69
 - first impressions, 66–67
 - games, 163
 - menus
 - adding all menus to Panel, 154
 - adding Applications menu to Panel, 154–155
 - deleting, 155
 - overview, 154
 - overview, 66, 152
 - panels, personalizing, 156
 - quick desktop guides, 69–71
- /dev directory, 219, 230
- devbootstrap directory, 230
- devel libraries, GTK+ 1.2, 432
- /dev/hda primary master, 460
- /dev/hdc secondary master, 460
- Device Manager, 79, 88
- df command, 238, 518
- DHCP (Dynamic Host Configuration Protocol), 37, 82
- DHCP (Dynamic Host Control Protocol), 82
- dialout system group, 440
- dial-up Internet access, configuring, 90–92
- Dictionary, 177–178
- diff command, 254, 518
- diff3 command, 254, 518
- dig command, 518
- digital cameras, transferring photos from
 - overview, 111
 - using camera's cable connection, 113
 - using card reader, 111–113
- digital images, 30
- digital music
 - creating own CDs, 288–289
 - legalities of playback, 279–280
 - listening to audio CDs, 283–284
 - online radio stations, 284
 - overview, 279
 - playing music files
 - installing codecs, 280–281
 - overview, 280
 - using Rhythmbox, 282–283
 - ripping music from CDs
 - adding MP3 support to Sound Juicer, 286–287
 - choosing format, 285–286
 - overview, 285
 - process, 287–288
- Digital Rights Management (DRM), 279, 296
- Dillo web browser, 431
- dip system group, 440
- Direct Memory Access (DMA), 455
- directing output. *See* piping and directing output
- directories, 185, 217

- Disk Cleanup tool, 28
 - disk drive, adding another
 - configuring Ubuntu to use, 461
 - overview, 460
 - partitioning disk, 460–461
 - Disk Management tool, 25
 - Disks Manager program, 193–194
 - Display Power Management section, 158
 - Distorts filters, GIMP, 313
 - distributions of Linux, 16
 - DivX format, 292
 - DMA (Direct Memory Access), 455
 - dmesg command, 519
 - DNS server addresses, 82
 - dns-clean service, 449
 - Do Backups drop-down list, Time tab, 467
 - Document area, OpenOffice.org, 330
 - Document protection, OpenOffice.org, 321
 - documentation. *See* manual
 - Documents category, OpenOffice.org, 332
 - dollar symbol (\$), 207
 - DoS (denial-of-service) attacks, 135
 - DOS (FAT), File System Type box, 198
 - DOS command prompt, 203
 - dosfsck * command, 519
 - dots per inch (DPI), 301
 - Download Package Files Only, 422
 - downloading Wine, 401–403
 - DPI (dots per inch), 301
 - dpkg command, 519
 - installing packages, 424–425
 - overview, 424
 - querying packages, 426
 - uninstalling packages, 426
 - dpkg -I packagename.i386.deb |less querying package, 426
 - dpkg -l packagename querying package, 426
 - dpkg-reconfigure command, 519
 - dpkg -s packagename |less querying package, 426
 - dragging-and-dropping files, 190
 - Drawing category, OpenOffice.org, 332
 - drawing programs, 162
 - Drawing toolbar
 - Evolution, 167
 - Impress, 366, 369
 - drive references, 183
 - drivername, 89
 - drivers, proprietary vs. open-source, 78–79
 - DRM (Digital Rights Management), 279, 296
 - DSL, 81, 98
 - du command, 237, 519
 - dual-boot setup, 25
 - dummy C:/ drive, 408
 - DVDs, 296–297
 - dw command, vi command mode, 245
 - Dynamic Host Configuration Protocol (DHCP), 37, 82
 - Dynamic Host Control Protocol (DHCP), 82
- E**
- e2fsprogs-udeb file, 50
 - echo command, 208
 - Edge Detect filters, GIMP, 312
 - Edit category, OpenOffice.org, 332
 - editing images. *See* image editing
 - Edubuntu, 538
 - Edutainment Suite, KDE, 538
 - Effect Options dialog box, 367
 - Eject button, 197
 - eject command, 519
 - ejecting media from drives, 197–198
 - Elliptical selection tool, GIMP, 303
 - Emacs text editor, 244, 269
 - e-mail
 - backing up, 30–31
 - composing messages, 388
 - configuring, 103–105
 - creating new folders, 390–391
 - creating search folders, 394

- creating signatures, 389–390
- deleting messages, 386–387
- flagging messages, 387–388
- junk e-mail, 391
- overview, 390
- reading, 386
- sending and receiving, 385–386
- sorting and filtering messages, 392–393
- E-mail option, OpenOffice.org, 338
- e-mail programs, 162, 168–169
- emblems, 188
- Emphasis option, Add button dialog box, 366
- Enable button, 193
- encoding, 280
- encryption key, 85, 481, 484
- End User License Agreements, Microsoft's, 404
- end-of-file (EOF), 270
- End-User License Agreement (EULA), 15
- Enhance filters, GIMP, 312
- entering data, in Calc, 355–356
- Entrance option, Add button dialog box, 366
- EOF (end-of-file), 270
- equal sign (=), Calc, 355
- Erase tool, GIMP, 305
- error dialog box, 406
- error messages, 49
- Escape key, Totem Movie Player, 173
- /etc directory, 230, 232
- /etc/crontab/ file, 473
- /etc/fstab configuration file, 406
- /etc/hdparm.conf file, 455
- /etc/inittab file, 211
- Ethernet cable, 108
- Ethernet card
 - configuring Network Card via DHCP, 82
 - configuring static IP address, 82–83
 - overview, 81
- Ethernet port, 93
- EULA (End-User License Agreement), 15
- evms service, 449
- Evolution (email/personal information management program), 19, 103, 168–169
 - Calendars mode
 - adding or editing diary entry, 397–398
 - overview, 396
 - specifying appointment types, 396
 - Contacts mode
 - adding or editing contact information, 395
 - creating contact list, 396
 - overview, 394
 - Mail mode
 - composing message, 388
 - creating e-mail signature, 389–390
 - creating new folders, 390–391
 - creating search folders, 394
 - deleting messages, 386–387
 - flagging messages, 387–388
 - junk e-mail, 391
 - overview, 390
 - reading e-mail, 386
 - sending and receiving e-mail, 385–386
 - sorting and filtering messages, 392–393
 - modes, overview of, 383–385
 - overview, 383
 - Tasks mode, 398–399
- ex * command, 519
- exclamation mark (!), 271, 484
- .exe file installer, 413
- Execute Actions drop-down list, Evolution, 393
- exit command, BASH, 519
- exit command, sftp, 485
- Exit option, Add button dialog box, 366
- Exit Saving Changes option, 34
- Expand box option, 156
- ext2 file system, 231

ext3 file system, 231

Extract button, Archive Manager, 177

extracting files from tar archive, 470

F

Face Browser feature, 151

FAT32 partitions, 192

fax system group, 441

fdformat command, 519

fdisk command, 233, 519

fetchmail service, 449

fg command, 519

fields, Writer, 352

file browser window, Nautilus, 188, 195

file command, 519

file compatibility, OpenOffice.org, 321–323

File Format drop-down list, 167, 370

file formats, OpenOffice.org, 341

file management

- accessing removable storage devices
 - ejecting media from drives, 197–198
 - formatting floppies, 198
 - overview, 196–197

accessing Windows files

- accessing networked files, 195–196
- overview, 192
- working with files in Windows partitions, 192–195

file system concepts

- case sensitivity, 184
- drive references, 183
- file access and storage, 184–186
- overview, 183

launching files and running programs, 190–191

overview, 183

using Nautilus

- file and folder icons, 188–189
- overview, 186–188
- special windows, 189–190

file permissions. *See* users and file permissions

file server, 538

file system, 228–231

file type codes, 224

File Types tab, 466

files. *See also* file management; users and file permissions

free space, 238

searching for

- overview, 235
- using find command, 235–236
- using locate command, 236–237
- using whereis command, 237

size of, viewing, 237–238

virtual files, 219–221

Files heading, man page, 531

fill screen option, Style drop-down list, 144

Fill tool, GIMP, 304

Filter function, Calc, 361

filters

- in Calc, 361
- in Evolution, 392–393
- in GIMP, 311–315

find command, 235–236, 520

Find Files function, Windows, 70

Firefox web browser, 169–170, 261

firefox-bin process, 261

Firestarter, 125

configuring

- overview, 132–133
- setting inbound rules, 133–134
- setting outbound rules, 134–135
- turning off diagnostic services, 135–136

installing, 131–132

firewalls, 131. *See also* Firestarter

FLAC audio file format, 286

flagging e-mail messages, 387–388

flags command-line, 214

Flash, 167, 363, 370

- flash animations, 297–299
- Flash file, exporting Impress presentation as, 370–371
- Flash Player software, 170, 297
- flatbed scanners, 114
- flavors of Linux, 16
- Flip tool, GIMP, 304
- floating palettes, OpenOffice.org, 329
- floating toolbar, OpenOffice.org, 340
- floppy disk drive, 196
- floppy disks, formatting, 198
- Floppy Formatter, 176–177
- floppy system group, 440
- flow charts, 320
- FluxBox desktop interface, 540
- folders
 - in Evolution's Mail mode, creating, 390–391
 - search folders, in Evolution, 394
- font preferences, 144–145
- Font Preferences dialog box, 144
- Font Rendering heading, Font Preferences dialog box, 144
- font-copying maneuver, 324
- fonts
 - copying Windows fonts, 323–325
 - installing TrueType Core Fonts, 324–326
 - overview, 323
- Fonts option, OpenOffice.org, 335
- Fontwork Gallery dialog box, 368
- Fontwork tool, 367–368
- Foomatic D1000 scanner, 534
- footers, in Writer, 352–353
- foreground (top) color box, 305
- Form Control toolbar, Base, 378, 380
- Form Navigation toolbar, Base, 381
- Form Wizard, Base, 377
- Format category, OpenOffice.org, 332
- Format Cell dialog box, Calc, 356
- formatting
 - data, in Calc, 355–356
 - floppies, 198
 - text, in Writer
 - context menu, 344
 - Formatting toolbar, 343
 - overview, 343
 - Style and Formatting palette, 345
- Formatting Aids option, OpenOffice.org Writer, 336–337
- Formatting toolbar
 - Calc, 355
 - Fontwork tool, 368
 - OpenOffice.org, 329
 - Writer, 343–344
- formatting tools, Writer, 164
- forms, Base, 167, 374
- Formula bar, Calc, 357
- formulas, Calc
 - overview, 356–357
 - summing figures, 358
 - using Function Wizard, 357
- forums, 533–534
- forward slash (/), 183, 530
- .4 command, vi command mode, 246
- FQDN (fully qualified domain name), 480
- Frame category, OpenOffice.org, 332
- framebuffer mode, 49, 55
- free command, 520
- Free Software Foundation (FSF), 13
- freeing disk space, 458
- fsck * command, 520
- FSF (Free Software Foundation), 13
- fstab configuration file, 232
- ftp * command, 520
- ftp program, 484
- full-screen option, 207
- fully qualified domain name (FQDN), 480

Function Wizard, Calc, 356–358
 functions, adding to toolbars,
 OpenOffice.org, 330–333
 fuser command, 520

G

g switch, 440
 G switch, 441
 Gaim, 105
 Gaim Internet messenger, 178
 GCC (GNU Compiler Collection), 14, 430
 GCompris, 538
 gcrontab program, 475
 gdm service, 449
 Gedit text editor, 119, 211, 445
 GeForce cards, 117
 General option
 Calc, 337
 Draw, 338
 Impress, 338
 Writer, 336
 General tab, GNOME Display Manager, 149
 Generic filters, GIMP, 313
 Geometry option, 3D effect, 369
 get command, sftp, 484
 GID (group ID), 439
 GIDs, 437
 GIMP (image editing program), 174–175
 cropping and cloning, 309–310
 filters, 311–315
 making color corrections, 308–309
 overview, 302–306, 308
 sharpening, 310–311
 Glass Effects filters, GIMP, 313
 glib libraries, GTK+ 1.2, 432
 Global Regular Expression Print, 251
 Globespan chipset, 92
 GNOME Art web site, 142–143
 GNOME CD/DVD Creator, 464
 GNOME desktop, running remotely, 486–487
 GNOME Display Manager, 149
 GNOME Onscreen Keyboard (GOK), 150
 GNOME Terminal, 94, 102, 119, 152, 207, 221,
 264, 270, 293, 447
 GNOME themes, 140
 gnome-app-install package, 424
 GnomeMeeting, 180
 gnome-session process, 487
 Gnopernicus, 150
 GNU and Linux together, 15–16
 GNU C Library, 416
 GNU Compiler Collection (GCC), 14, 430
 GNU organization, 5
 GNU project, 13–14
 GNU Public License (GPL), 14–15, 90
 Gnumeric, 326
 Go Back option, 39
 gocr program, 114
 GOK (GNOME Onscreen Keyboard), 150
 Google Summer of Code sponsorship
 scheme, 464
 GPL (GNU Public License), 14–15, 90
 Gradient fill tool, GIMP, 305
 Graphic category, OpenOffice.org, 332
 graphical boot process, 44, 51
 graphical problems
 graphics card configuration, 54–55
 keyboard and mouse configuration, 55–57
 monitor configuration, 57–59
 overview, 53–54
 graphical subsystem, 259
 graphical user interface. *See* GUI
 graphics. *See* image editing
 grep command, 274, 520
 Grid option
 Calc, 337
 Draw, 338
 Impress, 338
 Writer, 336–337
 group ID (GID), 439

- groupadd command, 440, 520
- groupdel command, 441, 520
- groups, creating and deleting via GUI, 439
- groups command, 520
- Groups tab, Users and Groups program window, 439
- GRUB boot loader, 42–43, 443
- GRUB boot menu, 52
- GRUB> command prompt, 52
- gstreamer packages, 281
- GSynaptics, 102
- gThumb software, 113, 301
- GTK+ Greeter tab, GNOME Display Manager, 149
- gtk2-engines search term, 142
- GUI (graphical user interface), 16
 - adding and deleting users via, 437–439
 - creating and deleting groups via, 439
- GUI applications, accessing remotely
 - accessing Ubuntu via Remote Desktop, 488–489
 - overview, 485
 - running GNOME desktop, 486–487
 - running X applications on remote computer, 485–486
- GUI terminal emulator, 270
- gzip command, 520
- gzip program, 469

H

- h command option, 237
- hackers, 127
- halt command, 520
- halt service, 449
- Hancom Office, 326
- Hand-drawn selection tool, GIMP, 303
- Happy GNOME with Browser option, 151
- hard disk
 - erasing, 39
 - settings, optimizing
 - benchmarking hard disk, 452–453

- changing disk settings, 453–455
 - making disk optimizations permanent, 455–456
 - overview, 452
 - space, freeing up, 28
 - spinning down, 158–159
- hardware
 - overview, 77–78
 - proprietary vs. open-source drivers, 78–79
 - viewing, 79–81
- hash (#), 208
- hdparm command string, 455, 521
- hdparm service, 449, 452
- hdparm settings, 455
- head command, 242–243, 521
- head mytextfile command, 242
- headers, Writer, 352–353
- help
 - online
 - forums, 533–534
 - mailing lists, 534–535
 - other official sites, 535
 - overview, 533
 - third-party sites, 535
 - in OpenOffice.org, 339
 - overview, 529
 - via manual
 - info pages, 532
 - man pages, 529–532
 - overview, 529
 - README files and other documentation, 533
- Help Agent, 339
- help command option, 268
- help command, BASH, 521
- help command, sftp, 485
- :help command, vi command-line mode, 248
- help option, 209
- hexadecimal (hex) keys, 86

- hidden files and directories, 190
- hierarchical file system, 228
- high-resolution mouse, 146
- history command, 521
- history of Linux, 11–13
- /home directory, 127, 222, 228, 230, 233, 463
- /home/ folder, 466
- Home Folder option, 148
- host command, 521
- hostname
 - command, 521
 - setting, 37
- hotkey-se, 449
- hotplug service, 449
- hotplug-n, 450
- HowStuffWorks site, 92
- HP JetDirect, 109
- hplip service, 450
- HTML file types, 191
- Hue/Saturation option, GIMP, 308
- Hurd, 16
- hyperlinks, 532
-
- i command option, grep, 251
- i command, vi command mode, 246
- .i386.deb file extensions, 426
- icon selection dialog box, 284
- icons, 141
 - for files and folders, 188–189
 - Ubuntu desktop, 68
- id command, 440, 521
- IDE1 slave (hdb), 39
- ifconfig command, 209, 521
- ifrename service, 450
- ifupdown service, 450
- Illumination option, 3D effect, 369
- ILS (Internet Locator Service) directory, 180
- image editing
 - getting pictures onto PC, 301
 - GIMP
 - cropping and cloning, 309–310
 - filters, applying, 311–315
 - making color corrections, 308–309
 - overview, 302–308
 - sharpening, 310–311
 - overview, 301
- image editor programs, 162
- Image file types, 191
- Image-Editing Tools, GIMP, 303
- image-tweaking tools, Writer, 346
- IMAP mail servers, 103
- Impress (presentation program), 166–167
 - animating slides, 366–367
 - applying Fontwork, 367–368
 - creating quick presentation, 363–364
 - exporting presentation as Flash file, 370–371
 - overview, 363
 - using 3D effects, 369–370
 - working in, 365–366
- inbound data connection, 132
- inbound traffic policy, 133–134
- Include Threads drop-down box, Evolution, 394
- incremental backup, 466
- industrial-strength operating system, 7
- info command, 521
- info pages, 532
- init command, 521
- initialization and run-level services, optimizing
 - disabling run-level scripts, 447–451
 - overview, 446–447
 - repairing problems, 451–452
- initialization scripts, 444
- initrd.img -> boot/initrd.img-2.6.12-9-386 directory, 230
- Ink tool, GIMP, 305

- input devices, personalizing
 - keyboard settings, 147–148
 - mouse options, 146–147
 - overview, 146
- input/output (I/O) mode, 454
- Insert category, OpenOffice.org, 332
- Insert mode, vi text editor, 246–247
- inserting data and cells, in Calc, 356
- Install Theme button, Theme Preferences dialog box, 143
- installing
 - compiler tools, 430–431
 - Firestarter, 131–132
 - GRUB boot loader, 42–43
 - Microsoft Office, 406–407
 - problems with, 47–51, 413. *See also*
 - graphical problems; postinstallation problems; preinstallation problems
 - RealPlayer, 293
 - theme components
 - downloading from GNOME Art web site, 142–143
 - downloading from Ubuntu repositories, 142
 - overview, 142
 - TrueType Core Fonts, 324–326
 - Ubuntu (overview), 33–34
 - Ubuntu (steps in order)
 - step 1-- booting from CD-ROM, 34–35
 - step 2-- selecting language, 36
 - step 3-- confirming keyboard layout, 36
 - step 4-- waiting during hardware and networking autodiscovery, 37
 - step 5-- setting hostname, 37
 - step 6-- partitioning disk, 38–40
 - step 7-- waiting during automated package copying, 40–41
 - step 8-- setting time zone, 41
 - step 9-- setting real name, username, and password, 41–42
 - step 10-- installing GRUB boot loader, 42–43
 - step 11-- rebooting for second installation phase, 43
 - step 12-- waiting during package installation, 44
 - step 13-- configuring monitor, 44–45
 - step 14-- booting for first time, 45–46
- installing software, 129
 - APT and package repositories
 - categories of repositories, 418–419
 - overview, 417–418
 - repository sections, 419–420
 - dependency management, 416–417
 - formats of Linux installation files, 414–415
 - overview, 413
 - package management from command prompt
 - overview, 424
 - using APT tools, 426–430
 - using dpkg, 424–426
 - packages, 415–416
 - from source
 - compiling, 432–434
 - installing compiler tools, 430–431
 - overview, 430
 - unpacking tarball and solving dependencies, 431–432
- Synaptic Package Manager
 - installing software, 422–423
 - overview, 420
 - removing software, 424
 - searching for software, 420–422
- instant messaging, setting up, 105–106
- Internet
 - configuring dial-up access, 90–92
 - overview, 81
 - USB-based DSL modem
 - Alcatel SpeedTouch series modem setup, 95–98
 - BT Voyager modem setup, 93–95
 - overview, 92

- using Ethernet card
 - configuring Network Card via DHCP, 82
 - configuring static IP address, 82–83
 - overview, 81
- wireless networks
 - configuring WEP, 86
 - configuring wireless card, 84–85
 - configuring WPA, 87–90
 - overview, 84
 - using WEP or WPA encryption, 85–86
- working with proxy server, 98–99
- Internet Explorer favorites, 30
- Internet Locator Service (ILS) directory, 180
- IO_support setting, 454
- IP address, 82–83
- iptables, 125
- ISO file, 50
- ISP login, 97

J

- j option, 470
- Java option, OpenOffice.org, 335
- Java Runtime program, 299
- jfs file system, 231
- jobs, controlling, 264–265
- jobs command, 521
- junk e-mail, 391
- junk mail filtering, Evolution, 391

K

- K Desktop Environment project (KDE), 537
- keepsettings setting, 455
- kernel, 4, 12, 443
- kernel files, 77
- Kernel Panic errors, 48
- kernel program, 6
- kernel-tree, 458
- key combination, 148

keyboard

- configuration, 55–57
- layout, 36, 56
- settings, 147–148
- shortcuts, 148
 - in BASH, 269
 - overview, 269
 - for system control, 270–271
 - for working in BASH, 269–270
- using Bluetooth keyboard, 122–123

- Keyboard Preferences dialog box, 147
- keyboard-based navigation system, 36
- kill command, 521
- killall command, 263, 521
- klogd service, 450
- KMail program, KDE, 537
- Knoppix, 29
- KOffice, 326
- Konqueror, 423
- Kubuntu, 537
- kubuntu-desktop package, 537
- KVM switch, 50

L

- l option, 221
- LAMP acronym, 20
- language, selecting, 36
- Languages settings, 336
- last command, 521
- launchers, 152
- launching files and running programs, 190–191
- layers, GIMP, 307
- Layout Options tab, Keyboard Preferences dialog box, 147
- Layouts tab, Keyboard Preferences dialog box, 147
- lcd command, sftp, 484
- ldd command, 521
- learning to use Linux, 19–20

- left angle bracket (<), 254
 - Left/right cursor key shortcut, BASH, 269
 - less command, 242, 522
 - lib directory, 230
 - libglib package, 432
 - libgtk package, 432
 - libraries, 203, 456
 - Light Effects filters, GIMP, 313
 - Linmodems, 90
 - Linux
 - age of, 5–6
 - benefits of, 8–9
 - flavors of, 16
 - learning to use, 19–20
 - obtaining, 21
 - what it is, 3–5
 - who uses, 20
 - Linux kernel, 219
 - Linux swap-ins, 19
 - linux-res, 450
 - listing files, BASH shell, 213–214
 - live distros, 540
 - lls command, sftp, 484
 - mkdir command, sftp, 484
 - ln command, 522
 - local printer, 107–108
 - locate command, 236–237
 - Location bar, Nautilus window, 187
 - logging in, 64–65
 - logging in to remote computer, 480–482
 - login options, 149–150
 - login picture, 151
 - long option (-l), 214
 - lost+found directory, 230
 - lpadmin system group, 441
 - lpr command, 522
 - lpstat command, 522
 - ls command options, 214, 522
 - ls command, sftp, 484
 - ls -l command option, 223, 237
 - ls of command, 522
 - lsmod command, 522
 - lvm service, 450
- M**
- m command option, 439
 - m hdparm command option, 454
 - Macromedia Flash, 167, 363, 370
 - macros, recording, OpenOffice.org, 340
 - Magnifier, 150
 - Mail Merge E-Mail option, OpenOffice.org Writer, 337
 - Mail Merge Wizard, Writer, 349
 - mail merging, Writer, 349–352
 - Mail mode, Evolution, 384
 - composing messages, 388
 - creating e-mail signature, 389–390
 - creating new folders, 390–391
 - creating search folders, 394
 - deleting messages, 386–387
 - flagging messages, 387–388
 - junk e-mail, 391
 - overview, 390
 - reading e-mail, 386
 - sending and receiving e-mail, 385–386
 - sorting and filtering messages, 392–393
 - Mail Preferences icon, Preferences dialog box, Evolution, 386
 - mailbox (.mbox) files, Thunderbird's, 388
 - mailing lists, 534–535
 - Main Distribution repository, 418
 - Main section, 419
 - main toolbar window, GIMP, 302
 - Main work area, 365
 - MainActor application, 299
 - make command, 433
 - make install command, 433
 - Make Link option, 153

- makedev service, 450
- makefile command, 433
- man command, 522
- managing users. *See* users, managing
- manual
 - info pages, 532
 - man pages, 529–532
 - overview, 529
 - README files and other documentation, 533
- Manual backups mode, 466
- Map filters, GIMP, 314
- Mark for Complete Removal option, 424
- Mark for Removal option, 424
- Mark Recommended for Installation option, 422
- Mark Suggested for Installation option, 422
- Massachusetts Institute of Technology (MIT), 14
- Material option, 3D effect, 369
- Max Size tab, 466
- MCU (Multipoint Control Unit), 180
- md5sum command, 522
- mdadm service, 450
- mdadm-raid service, 450
- Measurer tool, GIMP, 304
- media directory, 230
- membership to Linux community, 9
- memory cards, 301
- Memory option, OpenOffice.org, 335
- memory stick, USB, 116
- Menu bar
 - Nautilus window, 186
 - OpenOffice.org, 329
- menus
 - adding all menus to Panel, 154
 - adding Applications menu to Panel, 154–155
 - customizing, in OpenOffice.org, 333–334
 - deleting, 155
 - overview, 154
- Menus element, 68
- metapackages, 423
- Microsoft, 3
- Microsoft Access, 167
- Microsoft Exchange, 103
- Microsoft Office
 - running under Ubuntu
 - installing Office, 406–407
 - installing other Windows programs, 409
 - overview, 401
 - running Office components, 407–409
 - Wine, 401–406
 - similarities with OpenOffice.org, 319–320
- middle mouse button, 74
- Mines, 179
- Minix, 11
- Mist theme, 140
- MIT (Massachusetts Institute of Technology), 14
- MIT Artificial Intelligence Lab, 15
- mkdir command, 217, 484
- mkfs * command, 522
- mkisofs command, 471, 523, 532
- /mnt directory, 230, 233
- /mntl/windows directory, 234
- Modem Port drop-down list, 91
- modems, USB-based DSL
 - Alcatel SpeedTouch series modem setup, 95–98
 - BT Voyager modem setup, 93–95
 - overview, 92
- modern Linux development, 17
- Modify category, OpenOffice.org, 332
- modinfo command, 523
- modprobe command, 523
- module-in, 450
- monitor configuration, 44–45, 57–59
- monitor resolutions, 59
- monitor shutoff, controlling, 157–158
- more command, 241, 523

Motion path option, Add button dialog box, 366
 Motion Picture Association of America (MPAA), 296
 Motion tab, Mouse Preferences dialog box, 146
 mount command, 220, 523
 mount point, 194
 mounting, 183

- drive, manually, 233–234
- overview, 232
- removing mounted system, 234–235
- using mount command, 232

 mountvirt, 450
 mouse, 73–74

- Bluetooth, 122–123
- configuration, 55–57
- options, 146–147

 Mouse Preferences dialog box, 146
 Move tool, GIMP, 304
 movie playback programs, 172–173
 movie/DVD players, 162
 movie-editing software, Linux, 299
 moving files, BASH shell, 215–216
 Mozilla Foundation, 30
 Mozilla Plug-in option, OpenOffice.org, 338
 Mozilla Thunderbird e-mail client, 30
 MP3 audio file format, 286
 MP3 players, 162
 MP3 support, adding to Sound Juicer, 286–287
 MPAA (Motion Picture Association of America), 296
 ms-dos prompt, 162
 multcount setting, 454
 multimedia

- DVDs, 296–297
- flash animations, 297–299
- overview, 291

TV

- checking for video input, 299
- installing tvtime, 300
- overview, 299

 videos

- installing codecs, 292
- installing RealPlayer, 293
- overview, 291–292
- playing back video, 293–295

 multimedia files, 30
 Multimedia Systems Selector, 299
 multiple virtual desktops, 139
 multiple-speaker surround sound, 123
 Multipoint Control Unit (MCU), 180
 Multiverse section, 420
 music. *See* digital music
 mutt e-mail program, 422
 mv command, 215, 523
 My Computer function, Windows, 70
 My Documents function, Windows, 70
 My Network Places, 195
 My Network Places/Network Neighborhood function, Windows, 70

N

n command option, 242
 N command, vi command mode, 246
 n command, vi command mode, 246
 name, setting, 41–42
 Name heading, man page, 530
 nano command-line text editor, 451
 nano text editor, GNU, 474
 Nautilus file browser, 112, 116, 140, 163, 173, 294, 463

- file and folder icons, 188–189
- overview, 186–188
- special windows, 189–190

 Navigate category, OpenOffice.org, 332
 ndd command, vi command mode, 245

- NdisWrapper, 90
 - netstat * command, 523
 - network cards, configuring via DHCP, 82
 - Network Name (ESSID) drop-down list, 89
 - Network Name (ESSID) field, 84
 - Network Neighborhood, 195
 - network printer configuration, 108–109
 - Network Settings applet, 86, 95, 98
 - networked files, accessing, 195–196
 - networking service, 450
 - New button, Evolution, 168
 - New button, Synaptic Package Manager toolbar, 421
 - New Panel option, 156
 - newbies, 9
 - Next Record button, 381
 - Next Track button, 283
 - NI column, top program, 259
 - nice command, 524
 - nohup command, 524
 - Noise filters, GIMP, 312
 - Nokia 6680 mobile phone, 119
 - Notes heading, man page, 531
 - Notification area, Ubuntu desktop, 68
 - ntpdate service, 450
 - nUbuntu, 540
 - Numbering category, OpenOffice.org, 332
 - Nvidia 3D cards, 116
 - Nvidia driver, installing, 117–118
 - nyy command, vi command mode, 246
- O**
- O command, vi command mode, 246
 - o command, vi command mode, 246
 - o option, 532
 - object modes, Nautilus, 189
 - obtaining Linux, 21
 - ODF (OpenDocument Format), 323, 341
 - Office, Microsoft. *See* Microsoft Office
 - Office file format, 321
 - Office files, 30
 - official sites, 535
 - Ogg media format, 280
 - Ogg Theora, 295
 - Ogg Vorbis audio file format, 286
 - one-off tasks, 477–478
 - online help
 - forums, 533–534
 - mailing lists, 534–535
 - other official sites, 535
 - overview, 533
 - third-party sites, 535
 - online software repositories, 99–102
 - online system upgrade, 52
 - online updates, 130–131
 - OpenDocument Format (ODF), 323, 341
 - OpenOffice.org, 373. *See also* Base; Calc; Spreadsheet; Writer
 - configuration options, 335
 - creating macros, 340
 - customizing interface
 - adding functions to toolbars, 330–333
 - adding new toolbar, 333
 - configuring options, 334–338
 - customizing menus, 333–334
 - overview, 330
 - file compatibility, 321–323
 - fonts
 - copying Windows fonts, 323–325
 - installing TrueType Core Fonts, 324–326
 - overview, 323
 - getting help, 339
 - inserting objects with Object Linking and Embedding, 339
 - interface, 329–330
 - key features, 320
 - overview, 319, 329
 - saving files, 341
 - similarities with Microsoft Office, 319–320

- toolbar, customization categories, 332
 - using wizards, 339
 - OpenOffice.org Basic, 165
 - OpenOffice.org BASIC Macros category, 332
 - open-source software, 322
 - open-source vs. proprietary drivers, 78–79
 - Opera, 163
 - operating system, defined, 4
 - opt directory, 230
 - Optimal Page Wrap, Writer, 347
 - optimization
 - adding another disk drive
 - configuring Ubuntu to use drive, 461
 - overview, 460
 - partitioning disk, 460–461
 - booting
 - initialization and run-level services, 446–452
 - overview, 443–445
 - reducing boot menu delay, 445
 - freeing disk space
 - emptying cache of package files, 459–460
 - emptying /tmp folder, 458–459
 - overview, 458
 - removing unused software, 460
 - hard disk settings
 - benchmarking hard disk, 452–453
 - changing, 453–455
 - making disk optimizations permanent, 455–456
 - overview, 452
 - overview, 443
 - prelinking
 - configuring prelinking, 457
 - deactivating prelinking, 457–458
 - overview, 456–457
 - options, configuring, in OpenOffice.org, 334–338
 - Options category, OpenOffice.org, 332
 - Options configuration dialog box, Writer, 352
 - Options heading, man page, 531
 - Optra, 107
 - outbound traffic, 133
 - outgoing connections, 134
- P**
- p command, vi command mode, 246
 - package conflict resolution, 417
 - package files, 414, 459–460
 - package management from command prompt
 - overview, 424
 - using APT tools
 - installing and removing packages, 427–428
 - overview, 426–427
 - querying packages and repositories, 429–430
 - updating system, 428
 - using dpkg
 - installing packages, 424–425
 - overview, 424
 - querying packages, 426
 - uninstalling packages, 426
 - package manager, 414
 - package repositories. *See also* APT and package repositories
 - package version, 281
 - package-manipulation command, 424
 - packages, 415–416
 - packaging, overview, 413
 - Paired Devices, 119
 - panel, Ubuntu desktop, 68
 - Panel entry, GIMP, 302
 - panels, personalizing, 156
 - Paragraph dialog box, Writer, 344
 - parent directory, 217
 - parent processes, 261

- partitioning, 25–27, 460–461
 - deleting Windows partition, 39
 - overview, 38
 - sharing hard disk with Windows, 38–39
 - using existing free space, 40
 - working with files in Windows partitions, 192–195
- passphrase, 86
- passwd command, 441, 524
- passwords, 128
 - adding and changing, 441–442
 - setting, 41–42
- pasting. *See* cutting and pasting text
- patents, software, 279
- Path creation tool, GIMP, 304
- Path field, 284
- PATH variable, 208
- paths, GIMP, 306
- Paths option, OpenOffice.org, 335
- Paths tab, 466
- Pattern stamp tool, GIMP, 305
- pax command, 468
- PCMCIA cards, 84, 102
- pcmcia service, 450
- PDF (Portable Document Format), 191, 320, 341
- PKSH (Public Domain Korn SHell), 204
- Pencil tool, GIMP, 305
- period (.), 235, 463
- permissions. *See also* users and file permissions
 - Permissions tab, 195
 - permissive policy, 133
 - personal ID number (PIN), 119
 - personal information management program. *See* Evolution
- personalization
 - desktop items
 - applets, 156
 - menus, 154–155
 - overview, 152
 - personalizing panels, 156
 - shortcuts, adding, 152–154
- font preferences, 144–145
- input devices
 - keyboard settings, 147–148
 - mouse options, 146–147
 - overview, 146
- keyboard shortcuts, 148
- login options, 149–150
- login picture, 151
- overview, 139
- power-saving features
 - controlling monitor shutoff, 157–158
 - overview, 157
 - spinning down hard disk, 158–159
- theme
 - changing individual theme components, 141–142
 - installing additional components, 142–143
 - overview, 139–140
 - wallpaper, 143–144
- Perspective tool, GIMP, 304
- Phone List view, Contacts mode, Evolution, 394
- photo editing program. *See* GIMP
- photos, transferring from digital camera
 - overview, 111
 - using camera’s cable connection, 113
 - using card reader, 111–113
- Photoshop 6, 409
- pictures. *See also* image editing
 - inserting, in Writer, 346–347
- PID column, top program, 258
- PID number, 261
- PIN (personal ID number), 119
- ping command, 524
- pipe symbol (|), 254

- pipng and directing output
 - of commands, 272–274
 - overview, 272
 - redirecting output, 274–275
- Places menu, 67
- Places pane, Nautilus window, 188
- plain text (passphrase) keys, 86
- plugdev system group, 441
- POP3 mail servers, 103
- Portable Document Format (PDF), 191, 320, 341
- postinstallation problems, 51–53
- Power Management box, 158
- powerowd service, 450
- PowerPoint, 166–167
- power-saving features, 49
 - controlling monitor shutoff, 157–158
 - overview, 157
 - spinning down hard disk, 158–159
- ppp service, 450
- pppd-dns service, 450
- PPPoA technical standard, 95
- PPPoE technical standard, 95
- PR column, top program, 258
- precompiled kernels, 458
- preinstallation problems, 47–49
- prelink command, 457
- prelinking, 456
 - configuring, 457
 - deactivating, 457–458
 - overview, 456–457
- prelinking scan, 457
- Presentation file types, 191
- presentation programs. *See* Impress
- Presentation Wizard, 167, 363
- primary key, Base, 379
- Print option
 - Calc, 337
 - Draw, 338
 - Impress, 338
 - Writer, 336–337
- Print Scr button, 32
- printenv command, 524
- printers
 - local printer configuration, 107–108
 - network printer configuration, 108–109
 - overview, 107
 - Windows/SMB shared printer configuration, 109–111
- proc directory, 230
- processes
 - controlling
 - killing processes, 260–262
 - overview, 260
 - using other commands for, 263–264
 - Zombie processes, 262–263
 - viewing, 257–260
- process-killing command, 263
- program window, Evolution, 383
- programs, 540. *See also* names of specific programs
 - audio playback, 171–172
 - CD/DVD burning, 173
 - database, 167
 - e-mail/personal information management, 168–169
 - movie playback, 172–173
 - overview, 161, 163, 175
 - photo editing, 174–175
 - presentation, 166–167
 - running, 71–72, 190–191
 - spreadsheet, 165–166
 - web browsing, 169–170
 - word processing, 163–164

Properties option, context menu, 188
 proprietary software, and GPL, 14–15
 proprietary vs. open-source drivers, 78–79
 Proxy option, OpenOffice.org, 338
 proxy servers, 98–99
 ps command, 263, 524
 Public Domain Korn SHell (PDKSH), 204
 put command, sftp, 485
 pwd command, sftp, 485, 524

Q

:q command, vi command-line mode, 248
 :q! command, vi command-line mode, 248
 queries, Database pane, Base, 375
 question mark (?), 242
 Quick Launch toolbar function, Windows, 70
 QuickTime format, 292

R

r command option, grep command, 251
 r command-line option, 215
 radio stations, online, 284
 RawWrite program, 47
 RDP (Remote Desktop Protocol), 489
 readahead service, 450, 455
 README files, 242, 533
 readonly setting, 455
 Real Video format, 292
 RealAudio streams, 284
 RealPlayer, 293
 reboot command, 524
 reboot service, 451
 receiving e-mail, 385–386
 Recommended backup mode, 465
 recording macros, OpenOffice.org, 340
 recovery mode, 53
 Rectangular selection tool, GIMP, 303
 Recurrence tab, Evolution, 397
 Recycle Bin function, Windows, 70
 redirecting output, 274–275

Regex tab, 466
 regexes, 252
 reiserfs file system, 231
 Reload button, 190, 420
 Remember Password box, Evolution, 385
 remote access to computers

- accessing GUI applications remotely
 - accessing Ubuntu via Remote Desktop, 488–489
 - overview, 485
 - running GNOME desktop, 486–487
 - running X applications on remote computer, 485–486
- connecting to remote Windows computers
 - other Windows computers, 491–492
 - overview, 489
 - Windows XP Professional, 2000, and NT, 489–491
- overview, 479
- using secure shell
 - logging in to remote computer, 480–482
 - overview, 479–480
 - transferring files between remote computers, 483–485

Remote Desktop, 488–489, 491
 Remote Desktop Protocol (RDP), 489
 remote shell session, 482
 Remote tab, 490
 removable storage devices

- ejecting media from drives, 197–198
- formatting floppies, 198
- overview, 196–197

removing

- unused software, 460
- Windows, 28

Render filters, GIMP, 314
 renice command, 524
 renice value, 260
 renicing, 260

- repartitioning hard disks, 29
 - Reply button, Evolution, 168
 - Reply To All button, Evolution, 168
 - reports, Database pane, Base, 375
 - RES column, top program, 259
 - resizing windows, 74
 - resolution, 67
 - Restore As button, 468
 - restoring data, via SBackup, 468
 - Restricted section, 419
 - restrictive policy, 133
 - reverse dependencies, 417, 426
 - Rhythmbox (audio playback program), 171–172, 280, 282–283
 - Rich Text Format (RTF), 241, 321
 - right angle bracket (>), 254
 - ripping music from CDs
 - adding MP3 support to Sound Juicer, 286–287
 - choosing format, 285–286
 - overview, 285
 - process, 287–288
 - rm command, 216, 485, 525
 - rmdir command, sftp, 485
 - rmmmod command, 524
 - rmnologin service, 451
 - root and ordinary users, 126–128
 - root directory, 230, 268
 - root password, 436
 - root user, 436–437
 - root user account, 126
 - Rotate tool, GIMP, 304
 - route * command, 524
 - router address, 82
 - rsync service, 451
 - RTF (Rich Text Format), 241, 321
 - Ruler, OpenOffice.org, 329
 - run level, 211
 - runlevel command, 524
 - Runlevel Control, 451
 - run-level scripts, 443
 - run-level services. *See* initialization and run-level services, optimizing
 - running programs, 71–72, 190–191
 - rw- file permissions, 226
- S**
- S column, top program, 259
 - s command option, 237
 - S command option, 237
 - Sanskrit, 206
 - SATA hard disks, 50
 - Save As dialog box
 - OpenOffice.org, 341
 - Writer, 164
 - Save As drop-down list, OpenOffice.org, 341
 - saving backup file to CD-R/RW, 471
 - SBackup, 464
 - changing backup file destination, 467
 - changing time period between backups, 467–468
 - excluding files and folders, 466
 - including files and folders, 466
 - overview, 464–466
 - restoring data via, 468
 - sbin directory, 230
 - Scale tool, GIMP, 304
 - scaled option, Style drop-down list, 144
 - scanner system group, 441
 - scanners, configuring, 114–116
 - scheduling tasks
 - with anacron, 476–477
 - with crontab
 - creating scheduled task, 474–475
 - editing with GUI, 475
 - overview, 473
 - overview, 473
 - using at command to schedule one-off tasks, 477–478
 - scp, 483

- screen * command, 525
- screen program, 482
- screen-cl, 451
- Screenreader, 150
- screensavers, 118
- script, defined, 403
- Search button, Synaptic Package Manager toolbar, 421
- search folders, in Evolution, 394
- Search option, OpenOffice.org, 338
- searching for files
 - overview, 235
 - using find command, 235–236
 - using locate command, 236–237
 - using whereis command, 237
- /second_disk directory, 461
- secure shell, remote access to computers using
 - logging in to remote computer, 480–482
 - overview, 479–480
 - transferring files between remote computers, 483–485
- Secure Shell (SSH) tool, 132
- Secure Socket Layer (SSL), 86
- security, 8
 - common-sense security, 128–129
 - Firestarter
 - configuring, 132–136
 - installing, 131–132
 - online updates, 130–131
 - overview, 125
 - root and ordinary users, 126–128
 - Ubuntu firewall, 131. *See also* Firestarter
 - Windows security vs. Linux security, 125–127
- Security option, OpenOffice.org, 335
- Security tab, GNOME Display Manager, 149
- Security Updates repository, 419
- Select Custom Icon option, 188
- self-installing binaries, 414
- Self-Installing Package, Windows, 491
- sending e-mail, 385–386
- sendsigs service, 451
- Sensitivity setting, 146
- Serpentine (CD/DVD burning program), 173, 288
- service, 446, 448–451
- Services Settings program, 447
- sftp, 484–485, 525
- Shading option, 3D effect, 369
- Shape selection tool, GIMP, 304
- sharing Linux, 8
- Sharpen filter, GIMP, 310
- sharpening, in GIMP, 310–311
- Shear tool, GIMP, 304
- shell commands, 272
- shell scripting, 275
- Shockwave Director browser plug-in, 299
- shortcuts, adding
 - creating launchers, 152–153
 - creating links, 153–154
 - overview, 152
- shortcuts, keyboard. *See* keyboard, shortcuts
- Shoutcast tab, 284
- Show Desktop icon, 154
- Show Hide Buttons, 156
- SHR column, top program, 259
- shred command, 525
- shriek, 271
- Shutdown/Reboot function, Windows, 71
- Sidenet script, 403, 407
- Sidenet Wine Configuration Utility, 403
- signatures, e-mail, 389–390
- single service, 451
- slash (/), 242
- sleep command, 525
- sliders, GIMP, 308
- slides, animating, 366–367
- Slides pane, 365
- smbclient * command, 525

- Smudge tool, GIMP, 305
- software, removing, 460. *See also*
 - installing software
- software download feature, Ubuntu, 324
- software patents, 279
- software toolkits, 126
- Software Updates window, 130
- Sophos antivirus labs, 126
- Sort By drop-down list, Calc, 359
- sort command, 275, 525
- Sort Lists option, OpenOffice.org Calc, 337
- sorting data, in Calc, 359
- sorting e-mail messages, 392–393
- Sound & Video submenu, 284
- sound cards, configuring, 123
- Sound Juicer, 286–287
- source code, 12, 414
- sources.list file, APT, 418
- space on hard drive. *See* hard disk, space, freeing up
- speed command option, 531
- spell-checking, in Writer, 345–346
- spin-down settings, 158
- spinning down hard disk, 158–159
- Spreadsheet file types, 191
- spreadsheet program. *See* Calc
- spreadsheets, 348
- square brackets, 252
- srv directory, 231
- SSH (Secure Shell) tool, 132
- ssh * command, 525
- ssh remote shell session, 480
- ssh server program, 480
- ssh suite, 479
- ssh-keygen program, 481
- SSID name, 89
- SSL (Secure Socket Layer), 86
- Stallman, Richard, 13–14
- standard error, 275
- standard input (stdin), 272, 275, 531
- standard output (stdout), 272, 275, 531
- Standard toolbar, OpenOffice.org, 329
- Star Office, 319
- Start menu function, Windows, 70
- starting up, 63–64
- startx command, 525
- static addresses, 31
- static IP address, 82–83
- status (S) column, 262
- Status bar, OpenOffice.org, 330
- stdin (standard input), 272, 275, 531
- stdout (standard output), 272, 275, 531
- stop-boot, 451
- storage, file, 184–186
- storage devices, removable
 - ejecting media from drives, 197–198
 - formatting floppies, 198
 - overview, 196–197
- storing backup data, 471
- Streamtuner, 284
- Style and Formatting palette, Writer, 345
- Stylus Color, 107
- su command, 525
- subdirectories, 185
- submenus, GIMP, 311
- subnet mask, 82
- sudo apt-get build-essential command, 431
- sudo apt-get dist-upgrade command-line, 428
- sudo apt-get update command, 428–429
- sudo command, 228, 437, 439, 525
- sudo passwd root command, 436
- sudo service, 451
- sudo system, 223
- Sum icon
 - Calc, 358
 - Writer, 348
- Summary field, Evolution, 397
- Sun Microsystems, 17, 319

- superusers, 126, 223
 - supplementary groups, 440
 - switches, command-line, 214
 - `:s/word/replacement/` command, vi
 - command-line mode, 248
 - Synaptic Package Manager, 21, 88, 99, 116, 142, 281, 284, 298, 300, 324, 326, 447, 457, 465, 475, 480
 - installing software, 422–423
 - overview, 420
 - removing software, 424
 - searching for software, 420–422
 - Synaptics TouchPad, 102
 - Synopsis heading, man page, 530–531
 - sys directory, 231
 - syslogd service, 451
 - system, controlling
 - controlling jobs, 264–265
 - controlling processes
 - controlling Zombie processes, 262–263
 - killing processes, 260–262
 - overview, 260
 - using other commands for, 263–264
 - overview, 257
 - viewing processes, 257–260
 - system control, keyboard shortcuts for, 270–271
 - system fonts, 144
 - system libraries, 5, 416
 - System menu, 67
 - System Restore, 28
 - system variable, 208
 - SysV Runlevel Config program, 447
- T**
- t command option, mount command, 234
 - t option, tar command, 470
 - Table category, OpenOffice.org, 332
 - Table option, Writer, 337
 - Table tool, Writer, 347
 - Table toolbar, Writer, 349
 - Table Wizard, Base, 375
 - tables
 - adding to databases, 375–377
 - Database pane, Base, 374
 - editing database table, 379–380
 - in Writer, 347–349
 - Tables option, OpenOffice.org Writer, 336
 - tac command, 525
 - tail command, 242–243, 525
 - tail mytextfile command, 242
 - Tanenbaum, Andrew, 12
 - tape system group, 441
 - tar archives
 - compressing, 469–470
 - extracting files from, 470
 - viewing information, 470
 - tar command, 477, 526
 - tar files, 468–469
 - tarball, unpacking, 431–432
 - .tar.gz archive, 143
 - task field, 477
 - Tasks, Evolution, 399
 - tasks, scheduling. *See* scheduling tasks
 - Tasks button, Evolution, 169
 - Tasks mode, Evolution, 385, 398–399
 - Tasks pane
 - Base, 374
 - Impress, 365
 - tee command, 526
 - templates, 363
 - Templates category, OpenOffice.org, 332
 - terminal emulator program, 207
 - terminal programs, 203
 - Terminal Server Client program, 489
 - Terminal window, GNOME, 102, 539
 - Texinfo system, 532
 - Text editor/viewer, 163

- text files
 - comparing, 254
 - overview, 241
 - searching through
 - overview, 251
 - using grep, 251–252
 - using regular expressions, 252–254
 - types, 191
 - vi text editor
 - creating new text file using, 250–251
 - modes, 244–248
 - overview, 244
 - using to edit file, 248–250
 - viewing
 - overview, 241
 - using cat command, 241
 - using head and tail commands, 242–243
 - using less command, 241–242
- Text Formatting toolbar, 167
- Text tool, GIMP, 304
- text-based shell, 204
- text-formatting tools, Writer, 164
- Textures option, 3D effect, 369
- Theme Details button, Theme Preferences dialog box, 141
- Theme Preferences dialog box, 140
- Themed Greeter tab, GNOME Display Manager, 149
- themes
 - changing individual components, 141–142
 - installing additional components
 - downloading from GNOME Art web site, 142–143
 - downloading from Ubuntu repositories, 142
 - overview, 142
 - overview, 139–140
- third-party partitioning tool, 40
- third-party sites, 535
- three dashes (---), 226
- three-button mouse, 57
- 3D effects, 369–370
- 3D flourish tool, 369
- 3D graphics card drivers, installing
 - ATI driver, 117
 - Nvidia driver, 117–118
 - overview, 116–117
 - testing 3D capabilities, 118
- Threshold setting, 146
- Thunderbird e-mail client, 30
- TightVNC system service, 491
- tilde (~), 474
- tiled option, Style drop-down list, 144
- Time bar, Totem Movie Player, 173, 294
- time zone, 41
- TIME+ column, top program, 259
- /tmp folder, 231, 458–459, 473
- Toolbar, Nautilus window, 186
- toolbars, in OpenOffice.org
 - adding functions to, 330–333
 - adding new toolbar, 333
- top * command, 526
- top program, 257
- Torvalds, Linus, 6, 11, 13
- Totem Movie Player, 172–173, 207, 291, 294
- touch command, 526
- Toys filters, GIMP, 315
- tpconfig utility, 102
- tracepath command, 526
- Track Editor button, 284
- transferring files
 - between Bluetooth devices
 - overview, 120
 - sending files from Ubuntu PC to another device, 121
 - sending files to Ubuntu PC, 120–121
 - photos from digital camera
 - overview, 111
 - using camera's cable connection, 113
 - using card reader, 111–113

- between remote computers
 - overview, 483
 - using scp, 483
 - using sftp, 484–485
- Trash folder, Evolution, 387
- TrueType Core Fonts, 323–326
- TV
 - checking for video input, 299
 - installing tvtime, 300
 - overview, 299
- TV card, 299
- TWAIN modules, Windows, 301
- Typing Break tab, Keyboard Preferences dialog box, 147

U

- u command, vi command mode, 246
- Ubuntu blog, 535
- Ubuntu Document Storage Facility, 535
- Ubuntu Guide, 535
- Ubuntu root file system, 230
- Ubuntu Server, 538–539
- ubuntu-desktop metapackage, 424
- udev service, 451
- udev-mtab service, 451
- UIDs, 437
- umask command, 526
- umount command, 234, 526
- umountfs service, 451
- unalias command, 526
- uname command, 526
- undo command option, 457
- Universe section, 419
- Unix operating system, 7, 11, 17
- unmaskirq setting, 454
- Unmount Volume, 113, 198
- unused software, removing, 460
- unzip command, 264, 526
- Update Manager icon, 130
- updates, online, 130–131
- Updates repository, 418
- uppercase letters, 184
- uptime command, 527
- urandom service, 451
- USB keyboard, 48
- USB memory stick, 93, 116
- USB socket, 219
- USB-based DSL modems
 - Alcatel SpeedTouch series modem setup, 95–98
 - BT Voyager modem setup, 93–95
 - overview, 92
- Use largest continuous free space option, 40
- User Account Editor dialog box, 437
- USER column, top program, 258
- User Data option, OpenOffice.org, 335
- User Privileges tab, 439
- useradd command, 439, 441, 526
- user-defined string, 274
- userdel command, 441, 526
- username, setting, 41–42
- Username box, 195
- users, managing
 - adding and changing passwords, 441–442
 - adding and deleting groups at command line, 439–441
 - adding and deleting users
 - at command line, 439–441
 - via GUI, 437–439
 - creating and deleting groups via GUI, 439
 - overview, 435
 - user and group accounts
 - overview, 435
 - root user, 436–437
 - UIDs and GIDs, 437
 - users and groups, 435–436
- users and file permissions
 - altering permissions, 227–228
 - overview, 221–223

- viewing permissions
 - overview, 223–225
 - permissions on directory owned by root, 226–227
 - permissions on user's directory, 226
 - typical data file permissions, 225–226
- Users tab, GNOME Display Manager, 150
- user-scheduled tasks, 473
- Using_Dma setting, 455
- usplash service, 451
- usr directory, 231
- /usr/share/doc directory, 533

V

- v option, 470
- var directory, 231
- /var/backup directory, 467
- variables, 208
- VBA (Visual Basic for Applications), 165, 321
- VBA compatibility, OpenOffice.org, 321
- vbesave service, 451
- vCard, Evolution, 396
- vector graphics creation tool, 320
- versions of Ubuntu
 - Edubuntu, 538
 - Kubuntu, 537
 - overview, 537
 - Ubuntu Server, 538–539
 - Xubuntu, 539–540
- vi * command, 527
- vi README command, 249
- vi text editor, 244
 - creating new text file using, 250–251
 - modes
 - Command mode, 245–246
 - Command-Line mode, 247–248
 - Insert mode, 246–247
 - overview, 244
 - overview, 244
 - using to edit file, 248–250

- video system group, 441
- videos
 - file types, 191
 - installing codecs, 292
 - installing RealPlayer, 293
 - overview, 291–292
 - playing back video, 293–295
- View As Icons/List, Nautilus window, 187
- View category, OpenOffice.org, 332
- View option
 - Calc, 337
 - Draw, 338
 - Impress, 338
 - Writer, 336–337
- VIRT column, top program, 259
- virtual desktops, 72–73, 139
- virtual files, 219–221
- Virtual Network Computing (VNC)
 - software, 488
- virtual PC software, 29
- viruses, 6, 126
- Visible Buttons entry, 330
- Visual Basic for Applications (VBA), 165, 321
- Vmlinuz -> boot/vmlinuz-2.6.12-0-386
 - directory, 231
- VMware, 7
- VNC (Virtual Network Computing) software, 488
- volume control applet, 123
- Voyager 205, 93
- vv option, 470

W

- w command, vi command mode, 246
- :w! command, vi command-line mode, 248
- :w command, vi command-line mode, 248
- wallpaper, personalizing, 143–144
- war driving, 85
- Wastebasket, 67, 155
- wc command, 527

- web browsers, 162, 169–170
- Web filters, GIMP, 314
- web page creation programs, 162
- web proxy, 98
- web server, 538
- websites, 535
- WEP (Wireless Equivalent Privacy), 85–86
- whatis command, 527
- whereis command, 237, 527
- widgets, 141
- Wi-Fi Protected Access (WPA), 85–90
- Winamp Lite audio software, 409
- Winamp media player, 413
- window borders, 141
- Window List, 154
- windows
 - Nautilus, special windows, 189–190
 - resizing, 74
- Windows computers, connecting to remotely
 - other Windows computers, 491–492
 - overview, 489
 - Windows XP Professional, 2000, and NT, 489–491
- Windows files, accessing
 - accessing networked files, 195–196
 - overview, 192
 - working with files in Windows partitions, 192–195
- Windows fonts, copying, 323–325
- Windows Media Player 9 format, 292
- Windows Start button, 154
- Windows XP Home machine, 324
- Windows operating system
 - deleting Windows partition, 39
 - problems with, 6–7
 - removing, 28
 - security of, vs. Linux security, 125–127
 - sharing hard disk with, 38–39
- Windows/SMB shared printer configuration, 109–111
- Wine
 - configuring, 403–406
 - downloading, 401–403
 - overview, 401
 - showing hidden files on CDs, 406
- .wine directory, 407
- WinZip archiving application, 409
- Wireless Connection, 84
- Wireless Equivalent Privacy (WEP), 85–86
- wireless networks
 - configuring WEP, 86
 - configuring wireless card, 84–85
 - configuring WPA, 87–90
 - overview, 84
 - using WEP or WPA encryption, 85–86
- wizard-based interface, 339
- wizards, in OpenOffice.org, 339
- word processing program. *See* Writer
- Word processor file types, 191
- Workspace Switcher tool, 72, 156
- worms, 6, 126
- WPA (Wi-Fi Protected Access), 85–90
- wpaupplicant software, 88
- write files, 193
- Writer (word processor), 163–164
 - adding headers and footers, 352–353
 - formatting text
 - context menu, 344
 - Formatting toolbar, 343
 - overview, 343
 - Style and Formatting palette, 345
 - inserting pictures, 346–347
 - mail merging, 349–352
 - overview, 343
 - spell-checking, 345–346
 - working with tables, 347–349
- Writer Options dialog box, 334
- Writing Aids settings, 336

X

- X command option, 455
- X flag, 485
- X graphical subsystem, 485
- x option, 470
- X server communications, 485
- x86-based processors, 430
- XDMCP tab, GNOME Display Manager, 149
- XFCE, 539
- xfs file system, 231
- xhost command, 527
- Ximian, 383
- xinit command, 527
- XMMS, configuring, 284
- X.org drivers, 116
- X.org graphical subsystem, 53
- xorg-comm, 451
- XSane utility, 114
- XServer tab, GNOME Display Manager, 149
- Xterm, 163
- xterm window, 487
- Xubuntu, 539
- XviD file format, 295

Y

- y option, 254
- yy command, vi command mode, 246

Z

- z command option, 470
- z option, 470
- zero (0) value, 445
- zip command, 527
- zip files, 162, 177
- zipgrep command, 527
- Zombie processes, 262–263
- Zoom controls, Nautilus window, 187
- Zoom tool, GIMP, 304, 306
- ZSH (Z SHell), 204

FIND IT FAST

with the **Apress *SuperIndex***[™]

Quickly Find Out What the Experts Know

Leading by innovation, Apress now offers you its ***SuperIndex***[™], a turbocharged companion to the fine index in this book. The Apress ***SuperIndex***[™] is a keyword and phrase-enabled search tool that lets you search through the entire Apress library. Powered by dtSearch[™], it delivers results instantly.

Instead of paging through a book or a PDF, you can electronically access the topic of your choice from a vast array of Apress titles. The Apress ***SuperIndex***[™] is the perfect tool to find critical snippets of code or an obscure reference. The Apress ***SuperIndex***[™] enables all users to harness essential information and data from the best minds in technology.

No registration is required, and the Apress ***SuperIndex***[™] is free to use.

- 1 Thorough and comprehensive searches of over 300 titles
- 2 No registration required
- 3 Instantaneous results
- 4 A single destination to find what you need
- 5 Engineered for speed and accuracy
- 6 Will spare your time, application, and anxiety level

Search now: <http://superindex.apress.com>

Super Index

Apress[®]
THE EXPERT'S VOICE[™]