

SUBSCRIBE or renew your subscription to APC for your chance to **WIN** the new Alfa 156 JTS, valued at over \$54,000



Only **\$65**
for 12 issues

apc
www.apcmag.com

Subscribe ...

■ Online at magshop.com.au or Call 13 61 16

Authorised under NSW Permit No. L02/09075 VIC: 02/2531 SA: T02/3553 ACT: TP02/3650 NT: NT02/3286
For terms and conditions refer to www.xmas.magshop.au. Expiry date: 24/12/02



THE **2003 edition linux**
POCKETBOOK

Editorial

Welcome back to *The Linux Pocketbook 2003* edition!

Many of you will probably remember the original print versions of *The Linux Pocketbook* on newsstands across the country. The original versions sold so well that we ran out of copies. We've had countless requests for reprints, so we've decided to bundle the entire book into this single resource. This version of the pocketbook relies heavily on Mandrake Linux 9.0 or Red Hat 8.0. Both were released late in 2002, and can be easily found for sale at www.everythinglinux.com.au, or for download from either mandrakelinux.com or redhat.com respectively. APC recommends that you use Mandrake Linux 9.0, which was bundled with the November issue of APC because it's a more technologically advanced distribution than Red Hat. This book is designed as a beginner's reference guide to the powerful (and initially confusing) world of Linux. Inside you'll find advice on partitioning and setting up your Linux system, tweaking, learning your way around the command line, working and even gaming under Linux. Release yourself from the shackles of commercial software. All you need is this pocketbook and a copy of Linux. The rest is up to you.

First published December 2000.
For more information about
The Linux Pocketbook 2003
edition call (02) 9288 9123.

Material contained within
The Linux Pocketbook 2003
edition is protected under the
Commonwealth Copyright
Act 1968. No material may be
reproduced in part or in whole
without the written consent of the
copyright holders.

The Linux Pocketbook 2003
edition is published by ACP Tech,
a division of ACP Publishing Pty
Ltd (ACN 053 273 546).

Matthew Overington

OBTAINING SOURCE

If you're a programmer who's interested in getting your hands on the complete source code for Red Hat 8.0 and Mandrake 9.0, the best place to go is directly to the source:

Red Hat 7.0: mirror.aarnet.edu.au/pub/linux/redhat/releases/redhat-8.0/SRPMS

Mandrake 7.2: mirror.aarnet.edu.au/pub/linux/mandrake/9.0/SPRMS

Alternatively, you can send a \$10 cheque or money order to the address below. Please include a return address along with the distribution source you would like, and we'll post you out a source CD.

Pocketbooks, GPO Box 4088, Sydney NSW 1028



Contents

CHAPTER 1	Customising Gnome	57	CHAPTER 6	WORKING WITH WINDOWS	131	
INTRODUCTION	11	Exploring KDE	60	What about Windows?	132	
The origins of the penguin	12	Customising KDE	64	Windows connectivity	138	
Why Linux?	18	CHAPTER 4	Basic security	145		
The ways of the world	20	USING LINUX	67	CHAPTER 7		
CHAPTER 2	Connecting to the Net	68	PLAYING WITH LINUX	151		
INSTALLING LINUX	21	Applications	71	Linux multimedia	152	
Choosing a distribution	22	Conjuring Linux	75	Multimedia playback	156	
Preparing to install Linux	23	Upgrading your system	84	Games	160	
Partitioning your drive	25	CHAPTER 5		CHAPTER 8		
Making space for Linux	28	LEARNING LINUX	85	LINUX RESOURCES	163	
Booting into Linux	32	Learning Linux	86	Getting help	164	
Installing Red Hat 8.0	33	Tools of the trade	95	Keeping in touch	166	
Installing Mandrake 9.0	40	Permissions	102	Uninstalling Linux	169	
CHAPTER 3	Shell operation	106	Mounting filesystems	110	Command reference	170
EXPLORING LINUX	47	Installing new applications	110			
Up and running	48	Compiling the kernel	117			
Understanding the desktop	51	Surviving in Linux	123			
Exploring Gnome	53	Configuration files	129			





Mandrake 9.0

A powerful installer, rich package list and stable operation help make the much-awaited sequel to Mandrake 8.2 a success for Linux old-hands and newbies alike.

Mandrake Linux has long been a favourite amongst both beginners and gurus, but version 8.2 suffered flak for having a poor installer and 'package bloat'. Not everyone can pick their way through installers that go bad halfway through, leaving a half-built filesystem and raw configuration files. Not every PC that runs Linux is able to store multiple versions of packages because the installer messed up the installation. In an emerging world of Linux distributions, complex, bloated installers are out. Package pileups are gone. Easy configuration and even easier operation is in. Mandrake 9 heralds a new wave of efficient Linux distributions that give Windows a run for its money in the usability stakes.

One of the sources of complaint regarding previous versions of Mandrake was the DrakX installer. DrakX was known as buggy and slow, and prone to installing multiple, conflicting versions of packages to the hard disk. DrakX has been refined and improved from previous versions, and is now one of the cleanest and most efficient Linux installers currently bundled with any distro. When we booted with the first disc in the drive, we were prompted to partition our disks. After setting up the file system, a menu is displayed where you can select the most appropriate packages for your PC. The package list for a distro like Mandrake is huge, so the list of programs is divided into appropriate categories for the tasks the computer might fill: Internet Appliance, Multimedia Workstation, and Server to name a few. From here the installer ran through cleanly, and installed a working version on each of six test machines without error. Totally painless.


IF LOOKS COULD KILL

We chose to install with both KDE 3.0.2 and GNOME 2.0 to give maximum flexibility. You can make both KDE 3.0.2 and GNOME 2.0 backwards compatible by installing the library files for previous versions (2.2 and 1.4 respectively), so you don't have to pick and choose which x-based applications you run. Galeon 1.2.5 is one of the standouts under either environment, as is the final release

version of Mozilla.

In case you've been looking to try one or more of the latest Linux applications, the best place to get them all in the one place is probably by installing Mandrake. Paging through the package list, Mandrake 9.0 reads like a who's-who of major Linux software development. The brand-new GNOME 2.0 is there, along with KDE 3.0.2 and version 2.4.19 RC1 of the kernel. In fact, all the packages included with Mandrake 9.0 are the latest and greatest so you won't need to spend much time, if any, downloading new packages and applying patches after installing it.

Mandrake 8.2 is one of the preferred Linux distributions on offer today, yet with release 9 Mandrake Linux appears to have improved on it. With GNOME 2.0, KDE 3.0.2, Xfree86 4.2, and a selection of new tools,



Mandrake 9.0 provides a clean look and feel through KDE 3.0.2.

Mandrake 9.0 is a must-have for any bleeding-edge Linux fan. But it's not just for gurus. Even the freshest Linux newbie will be able to install Mandrake 9.0 and be looking at a fully functional Linux desktop within an hour. Contact www.mandrakelinux.com for more information, or look to the October 2002 issue of *APC* for a working version.





Red Hat 8.0

Red Hat is one of the most popular Linux distributions on the market. As it speeds across the ether to PCs near you, APC asks what all the fuss is about.

Red Hat is striving to create a simple, powerful Linux distribution that can go head-to-head with Microsoft Windows in a corporate environment. The latest version comes with an office-focused set of productivity tools and a clean new look to try and woo corporate users towards Linux. Red Hat's size, relative success, and policy of creating an easy-to-use Linux distribution have caused many geeks to refer to it as the Microsoft of the Linux world. Striving for a unified appearance, simplified package installation system and robust driver set isn't going to hurt any Linux distribution though. Red Hat 8.0's core components include GCC 3.2, GTK+ 2, kernel 2.4.18, Gnome 2.0.1, KDE 3.0.3 and OpenOffice.org 1.0.1, keeping it in line with Mandrake 9.0. The only real difference is that Mandrake 9.0 comes bundled with kernel 2.4.19, which provides a few extra goodies, like improved FireWire handling and some updated drivers. Red Hat has stuck with version 2.4.18 since their 7.3 release, and has modded and patched it according to its own needs. This tends to make the kernel unnecessarily bulky, slow and hard to patch, but also gives the chance to improve security and add some minor bug fixes.

We downloaded the 5 ISO files from Red Hat's overloaded FTP site, and managed to get Red Hat 8.0 up and running within hours of its release. After booting from the first CD and being greeted by the now-familiar Anaconda installer, we were on our way in minutes. Anaconda has been updated slightly to include anti-aliased screen fonts and GTK+ 2, but otherwise it remains largely the same as previous versions. It picked up our Samsung SyncMaster monitor and Asus GeForce 3 card without any problems (although the GeForce 3 was detected as a 'generic GeForce 3').

The language support and driver collection have been beefed up slightly, and you are given 19 languages to choose from by default, including Ukrainian, Japanese and Korean. Our generic keyboard was picked up without a problem, but Anaconda was tripped up by our mouse. The mouse is not uncommon hardware — a USB Microsoft Intellimouse — but Red Hat only detected it as a Generic 3 button unit. We manually changed the setting to use the correct Microsoft Intellimouse driver, but we still didn't have use of the scroll wheel. It wasn't until we killed the X-server and

restarted it that we managed to get mouse support. Another bug. Given that the Microsoft Intellimouse is extremely popular, we'd have thought this was something Red Hat could get right (both Mandrake 9.0 and Lycoris Desktop/LX had no problems with the Intellimouse).


DRIVE SETUP

As per version 7.x, you can choose to autopartition your drives, launch Disk Druid (a GUI partitioning application) or take the full geek approach and use fdisk. Disk Druid has remained largely unchanged from Red Hat 7.3 and is still the preferred method for configuring your disks.

After the drives are set up, you are taken to the package selection screen. Red Hat has relied in the past on Gnome for its desktop environment, and Gnome 2.0 is selected as the default desktop under version 8.0, although KDE is also present. KDE is

not installed by default, but can be added with a few clicks when selecting the packages. Strangely, when you select KDE, a few packages, such as Kaboodle (the KDE media player) are disabled by default. KDE's relationship with Red Hat has been tenuous at best, culminating in KDE programmer Bernhard Rosenkraenzer quitting Red Hat during the development of version 8.0 over changes to KDE's

Red Hat 8.0 is disappointingly lacking in features.





code. KDE has been modified slightly in its appearance (and the annoying 'About KDE' box has been moved), but it retains its functionality.

You can manually select the packages to install, but we opted to go the whole hog and dump every package onto the hard disk. You'd be mad to opt for an installation as bloated as this, but it worked a treat for testing. The installer software took approximately an hour to complete, transferring 4.7GB worth of packages to the hard disk. During this time you are given updates on package progress, total progress, and a running tally of the packages already copied to the drive.

After everything was installed, we created a boot floppy and restarted the machine. The Red Hat Setup agent

kicked in and guided us through configuring the date and time, and then the configuration of the soundcard. If you register, you can connect to a network connection and run the update agent. This downloads and installs new packages much like the Windows Update feature does for Windows. If you refuse to register, you cannot access this feature.

Red Hat 8.0's Update Agent works much like Windows' Update feature to keep the OS up to date.

BLUECURVE

Bluecurve is an interesting innovation from Red Hat to try and force a unified look and feel across all applications. It's a new GTK+ theme that sits halfway between Gnome and KDE. Essentially it makes Gnome look more like KDE, and vice versa. Full versions of Gnome and KDE are there, and you can disable the Bluecurve interface if you wish. Bluecurve gives you some benefits like Antialiased text.

Unfortunately Red Hat's attempt at creating a single, unified theme has not quite hit the mark. Many programs that don't rely on the latest GTK+ still maintain their Gnome or KDE specific appearance. This goes for applications that rely on GTK+ 1, or KDE 1 or 2. This leads to a desktop environment where most windows and buttons look the same with a few glaring exceptions. Red

Hat is to be commended for attempting to unify the appearance of applications, as the dog's breakfast appearance of most X sessions is unappealing to people used to the single, unified experience of Windows, MacOS, or BeOS. Unfortunately all the applications should be updated to support a common theme or be left out. It sounds ruthless, but it's the only way to achieve the unified appearance that Red Hat is aiming for.

MENU SYSTEM

If Red Hat is aiming to get Red Hat 8.0 onto the desktop of corporate machines across Australia, then work has to be done on the menu systems. Both the Gnome and KDE menu systems are confusing, and it's hard to remember where to look to find a certain settings box. For example, there are four separate configuration sub menus if you are using Gnome, labelled 'Preferences', 'Server Settings', 'System Settings', and 'System Tools'. To the experienced user, this is bearable, but for some-

one that's new to Linux, this seems obscure, confusing and even alienating. If Red Hat is striving to create a simple, powerful distribution to gain acceptance in the corporate market, perhaps they should start by looking at Lycoris' menu screens. Lycoris Desktop/LX has a long way to go in terms of creating a solid installer and improving driver support, but at least the menu systems are heading in the right direction.

The package set that comes with Red Hat 8.0 is good: it comes with GCC 3.2, GTK+ 2, Gnome 2.0, and KDE 3.0.3. OpenOffice.org is included, as is KOffice 1.2, so you get a solid selection of the most common Linux applications. Red Hat has worked with OpenOffice to ensure that anti-aliased fonts are rendered correctly, which substantially improves the appearance when you're working on a text document. In keeping with its assault on the corporate environment, Red Hat 8.0 comes with Ximian Connector, which is useful for accessing Microsoft Exchange servers under Linux.





One of the best new additions is the inclusion of a dedicated fonts directory. If you create a folder called `~/fonts`, you can copy all your TTF fonts to that folder and they are accessible by most GTK+ applications. Unfortunately this doesn't extend to all GTK+ applications, and we found it frustrating that we'd lose access to our fonts under OpenOffice.org and Gedit. This sort of inconsistency is rife in Red Hat 8.0, which simply doesn't look polished. There seems to be a disturbing trend of releasing x.0 versions of software without adequate testing and completion of feature sets. This is not going to win over any users. Linux users are past the point of putting up with a kludgy operating system simply because it's free. It also has to work well.

The most surprising thing about Red Hat 8.0 is not what is installed, but rather what is left out. One thing that we couldn't believe after sitting through the installation process was that Red Hat 8.0 has virtually no multimedia support. Not even MP3 support. There is no Java, Flash, or Real media either. There is also no XINE, XMovie, or any player other than Kaboodle. Even in a distribution catering to the office user, we'd have expected to see some more advanced media options than are included by Red Hat by default. We're now all working in 2003, and media support in operating systems has been standard for half a decade. Irrespective of what you are using the computer for, you should be able to play MP3s, so Red Hat's business focus with version 8 is no excuse. Sure, you can always break the law by downloading and adding MP3 capabilities to your 8.0 installation, but why bother? It should be there in the first place. Thankfully, support for OGG and WAV files has been included by default, so Red Hat 8 is not completely devoid of audio support.

CLOSE... BUT NO CIGAR

Under the bonnet, GCC 3.2, GTK+ 2, and kernel 2.4.18 (although 2.4.19 could have been used), helps Red Hat to hum along. GCC 3.2 is noticeably faster than previous versions, and even though Red Hat is a bit bloated, it doesn't lag on modern hardware — it appears quicker than both MacOS X and Windows XP on identical hardware. All told, Red Hat 8.0 would not be too out of place as a business desktop environment, but Red Hat really has a lot to learn from modern distributions like Mandrake 9.0 or Debian 3.0r0 (Woody). Mandrake 9.0 is easy to install, replete with the latest packages, and stable. Debian can be customised and tweaked endlessly and provides a much stronger and more configurable operating platform. Red Hat's failure to include any sort of MP3 codec is inexcusable, even given the current arguments that are raging over DRM. Even though Bluecurve is a step in the right direction towards creating a unified desktop appearance and creating rules for the display of buttons, text and windows, it won't be effective until every application is updated to rely on GTK+ 2. In the meantime, you still end up with a cluttered, unappealing set of unmatched windows splayed all over your desktop. Our advice? Steer clear of Red Hat 8.0 and spend your time downloading Mandrake 9.0 instead.



Introduction

THE ORIGINS OF THE PENGUIN 12

WHY LINUX? 18

THE WAYS OF THE WORLD 20



The origins of the penguin What are the humble origins of this remarkable development in IT history, and where did it all begin?

Linux has grown rapidly during the past few years, and so has the IT world's awareness of it, following the media groundswell surrounding Linux in 1998. Back then it was the latest hot new thing to enter the realm of IT.

Now, the hype has levelled out and the facts are starting to establish themselves.

Developed under the Open Source model, Linux is built on reliable technology that performs well — and it's free. Its recent development has enabled it to establish a firm foothold in the world of computing. Linux has even taken Microsoft by surprise.

There isn't any real competition, however, as Linux isn't owned by any one person or entity. Moreover, Linux — by nature of its development — is continually improved upon only for the sake of itself. Linux's popularity is not the result of a marketing engine pushing the product around the world, but purely the result of users choosing something different, choosing to use an operating system that's designed to be the best it can be — an OS made by the people for the people.

THE BIRTH OF AN OPERATING SYSTEM

The roots of Linux trace back to the late 1960s and the American Telephone and Telegraph (AT&T) corporation. At the time, AT&T, in conjunction with the Massachusetts Institute of Technology (MIT), was working on an operating system called Multics. Multics had many problems; chief among them was the cost of running the system on the General Electric mainframe (GE 645). The development of the system was also proving to be unsatisfactory.

However, the group held on to Multics because it offered multiuser capabilities. The programmers could work together and share information easily, and they wanted to retain those freedoms at all costs. As time passed and financial pressures increased, the group began searching for a suitable replacement.

A member of the AT&T group, Ken Thompson, had begun designing

The roots of Linux trace back to the late 1960s and the American Telephone and Telegraph (AT&T) corporation.

a game called Space Travel, which simulated the movements of solar bodies. The user navigated through the celestial bodies and could land on the various objects. Unfortunately, the performance of the game was sluggish and cost too much to run. When Thompson found an unused Digital PDP-7 computer, he and colleague Dennis Ritchie rewrote the game in assembler and transferred it across using paper tape, at the same time implementing a filesystem that Thompson had been playing around with on paper.

During the process of porting the game, they also implemented a command interpreter and some basic utility functions for copying and removing files. In early 1970 another AT&T employee, Brian Kerningham,

1969

1969 AT&T moves away from Multics project, coinciding with the beginnings of Unix.



Linux timeline



1970

1970 Unix Epoch: Unix time begins counting on January 1.





suggested the name 'Unix' as a pun on Minix (a programming teaching tool). The foundations for a new operating system had been laid.

At this time, the Unix development team requested a PDP-11 machine from AT&T management. It was more powerful than the PDP-7, but less expensive than some of their previous requests. They managed to secure the system by proposing to develop a word processing and formatting machine. By the end of 1971, the PDP-11 was adopted by the AT&T patent office and Unix's reputation was beginning to build. Meanwhile, the team further developed the system.

Thompson had already written the B programming language, which began as an attempt to port Fortran and evolved into a new language. The C programming language was then developed as the successor to B, and by the end of 1973 the Unix kernel had been rewritten in C. At the time, AT&T wasn't allowed to sell the operating system, so it began distributing it to universities at a nominal cost. The combination of cheap hardware (the PDP-11) and cheap software (Unix) ensured the popularity of the system. The uptake of Unix in the university arena also ensured that hackers grew up using and adding to the system — a tradition that remains to this day.

A HACKER AND HIS FREEDOM

Around this time, a young hacker by the name of Richard Stallman began his career at MIT. He and his colleagues worked closely, sharing their code and their solutions to various problems. By the 1980s, however, most Unix variants had become proprietary, and its users could not view or modify the source code. For a while, Stallman hid from the encroaching threat of proprietary software in the artificial intelligence lab at MIT.

1973 Following on from the success of the newly created C language, the Unix operating system was completely rewritten in C. The C language was easily portable to other hardware and enticing to programmers, ensuring the system's popularity.

Linux timeline → → →

Stallman thought it was unethical to prevent people from sharing their ideas with each other. He wanted to live in a community where people were free to help each other, and free to improve upon their predecessors' works. After much thought and deliberation, he set out to create a completely free operating system. Free, in this sense, refers to the freedom to modify and distribute the source code — and not the price of the system. He decided to base this system on Unix, which had already been proven and well established. It would also allow him to test compatibility of the software, lessening the learning curve for existing Unix users.

Stallman named the system GNU, a recursive acronym for GNU's Not Unix. He began compiling existing free software and writing what remained. Once completed, he planned to release the system to anyone who chose to use it. The ultimate

Stallman thought it was unethical to prevent people sharing their ideas with each other.

goal was to provide enough free software (remember, we're talking about freedom, not price) so that anyone could use a computer without having to use proprietary software. His first project, the popular Emacs text editor, was completed in 1984 and released under the newly fashioned GNU Public License (GPL).

The GPL was designed to protect and promote the freedom of the user. The only restrictions it contains are those necessary to ensure that future users of the software will have the same freedom available to them. The licence states that it is acceptable to distribute or modify the source code — and even sell it — but that the source code must be distributed with the application, and any derived works must also be placed under the GPL.

1984 After deciding that it was unethical to use proprietary software, Richard Stallman began rewriting the Unix operating system from scratch. His first project, the Emacs text editor, was released under the GPL, which would become the most popular licensing scheme for Linux.

→ → →



1991: A KERNEL IS BORN

Linus Torvalds was a headstrong student at the University of Helsinki in 1991. Exposed to the power of Unix at university, but unable to afford the software or hardware, he began searching for a cheaper operating system that provided sufficient power. The operating systems of the day, however, left much to be desired. DOS didn't meet his needs, and the closest Unix replacement was the Minix teaching tool, developed by Andrew Tannenbaum.

A Unix-like operating system available for the Intel architecture, Minix proved to be very interesting to the young developer. Fortunately, it lacked many features that he required, so he was forced to do what any other right-minded hacker would have done — he started writing his own.

By October 1991, he had developed a workable kernel, which he named Freix — a cross between 'free Unix' and 'phreaks' (a hacker term). He managed to convince an associate to lend him some space on an FTP server so he could make the code available to others. The site's maintainer, however, didn't like the name Freix, and renamed it Linux as a play on Linus Torvalds' name.

When the 0.02 release of Linux was running competently on Intel hardware, Linus posted the following message to the comp.os.minix newsgroup:

"Do you pine for the nice days of Minix 1.1, when men were men and wrote their own device drivers? Are you without a nice project and just dying to cut your teeth on an OS you can modify for your needs?"

At first, a few people downloaded the code and began playing with it. Sometimes they sent in patches for inclusion in the kernel. Eventually, word of a freely available Unix clone spread, and what began as a small project grew

into an international operating system. Today, Linus estimates that only about 50% of the code in the kernel is his own.

The site's maintainer didn't like the name Friex, and renamed it Linux as a play on Linus Torvalds' name.

LINUX TODAY

Linux currently runs the majority of Web servers on the Internet, and is employed as a file and print server by thousands of corporations worldwide. The Linux kernel has proven to be extremely stable and versatile, and is packed with features.

and 20 million users worldwide. In addition, many commercial software vendors, including Corel, Informix, Oracle, Netscape and IBM are porting applications to Linux.

Previously, development was focused on enhancing support for symmetric multiprocessing (SMP), but the current focus for many developers is in the area of usability — bringing a coherent, integrated desktop environment into existence.

And now, you're holding it in your hand. Linux has been a long time developing and will still be developing years from now. Enjoy this unique time in computing history, as the future holds great promise.

Because Linux can be freely downloaded all over the world, estimating the current user base is almost impossible. However, estimated figures range between 10

You are entering the Linux community at a point in time where the system is entering a new phase. The distributions on the cover of this Pocketbook are cleaner, better integrated and easier to use than ever. Good luck, and may the source be with you.

1991 A young student studying computer science at the University of Helsinki begins a quest to find a suitable Unix-like OS for the Intel architecture. With limited choices available, Linus Torvalds begins writing his own, eventually releasing it on the Internet for public consumption.

Linux timeline → → →

2002 By late 2002, the Linux revolution is in full swing. Lycoris, Lindows, OEOne, Mandrake and Red Hat all release major new versions. Linux desktops are stable, feature-rich and replete with a wide-range of software packages. Several administrations announce plans to roll out Linux in the workplace, including the Chinese and Northern Territory governments.

← ← ←



Why Linux? It's clear that Linux just keeps on getting bigger and better. But why is Linux so popular and why do so many people make the switch to the penguin operating system?

What often piques the interest of new users is the price point of Linux — free. Building an operating system with the variety of features that Linux has is no small task, and while competing OSes with similar capabilities sell for hundreds or thousands of dollars, Linux doesn't cost a cent. You can download the entire operating system from the Web yourself, or buy it for nothing but the cost of the distribution medium. There aren't even any licence fees.

What also interests new users are the claims that Linux is more stable, efficient and feature-rich than other operating systems on the market. Sadly, Windows has a lot to answer for in the world of computing, because it has created a worldwide perception that computers are unreliable, and that's 'just the way things are'. But Linux is showing us that computers really can be relied upon, and are up to the task of running smoothly for months, or even years, without a hitch.

Invariably you will end up hearing about a new development model that is at the core of Linux — Open Source. The premise behind Open Source is the

public sharing of ideas and skills, of producing software with publicly available source code — a direct contrast to the closed development model that traditional software houses use.

The Open Source model allows software to be improved upon and added to by anyone, anywhere. As a result projects such as Linux have had thousands upon thousands of programming hours devoted to them by thousands of programmers around the world. A closed software development model simply can't compete with what Open Source can achieve in terms of reliability, performance and features.

Linux is the fastest growing operating system in the world today, and stands as a testament to the quality of open source software. (For more information visit www.opensource.org.)

Linux isn't the only popular Open Source project that makes news. According to Netcraft (www.netcraft.com), the Apache Web server runs approximately 60% of the world's sites, while its Microsoft competitor manages just over 20%. Sendmail, the Open Source mail transport agent, manages the majority of email

communication over the Internet. And these are just the more well-known examples.

The key to the Open Source movement is the GPL, or GNU Public License. Open source programs are not only distributed free of charge, but the magic building blocks of the software — the source code — is also distributed free. Anyone is welcome to make use of the source code, to add to it or modify it and create entirely new programs for their own purposes — including making money.

The GPL stipulates, however, that any changes made to the source code must be made available free of charge for others to do the same. In other words, you're welcome to benefit from the work of others as long as you allow others to benefit from yours. The availability of the source code encourages bug audits, modifications and improvements, and the licence ensures those modifications are distributed back to the general public.

From a business perspective this is a major advantage. If you're using Windows and a security flaw is found, you have to wait until Microsoft patches the hole, and in the meantime your business remains vulnerable. Under Open Source, if a security flaw is found you can set your own programmers, or employ freelance programmers, to patch it up immediately because you have access to the source. For the same reasons you can also modify and improve the program to suit how you want to do business, rather than learn to use software by a closed develop-

ment company that can't possibly understand the intricacies of how your business operates.

But what about the desktop? While Linux might not be quite as simple to use as Windows just yet, it's come far enough to almost dissolve the issue of 'ease of use'. The desktop environments provided in today's Linux distributions are not only simple to use but provide far more features than your humble Windows desktop. And you also have a choice as to which desktop environment you want to run. Linux can operate in many different ways, giving you the freedom to use your machine in whatever way is most comfortable to you.

Linux is also versatile. Out of the box you can configure Linux to be a fully functional desktop operating system, a secure and reliable server platform, a network router and firewall, a file and print server, and much more. Linux is an operating system you can mold to your own vision, to perform any number of tasks, with stability and performance, all for free. That's why so many people choose to use the popular penguin operating system.

If this is your first time looking into an alternative operating system, if you're wondering if there's a better way to use your computer, then this book has everything you need to get started with Linux and everything you need to help you realise the potential it offers.

Welcome to the 2003 edition of *The Linux Pocketbook*.

The ways of the world

Initially, Linux isn't going to be what you expect, and it will take time and patience on your part to learn the ropes.

First you have to be prepared to accept that Linux is likely to be foreign to you. Things won't work the Windows way in Linux, and for the most part, this is a good thing. Once upon a time you had to learn Windows from the ground up, and now you'll need to do the same with Linux. If you persevere, you will be rewarded with a rich and powerful operating system.

Second, Linux lets you live close to your system. You will encounter the bare bones of hardware and software with Linux, and how it works. Although this may be confronting at times, the advantage is that you will always have complete control of how your system works and behaves.

The more you understand about your machine, the better you will fare. It is possible to get by just using the graphical tools supplied with Linux. But in order to learn about your system and make the most of it, you'll need to look under the hood.

With this in mind, it needs to be said that *The Linux Pocketbook* 2003 edition is a concise guide that will get you started on your way with Linux, but it won't answer all your questions. The most powerful tool for any user is quality documentation, and Linux comes with a vast collection of guides and how-to documents to answer any question you may have.

If you find yourself stuck as you progress through this book, bear in mind that you're not the first and someone, somewhere has an answer for you. If you don't find it in the supporting documentation written by users just like you, you'll find it in World Wide Web (WWW) pages, IRC channels and Usenet groups, provided by users just like you. The Linux community is your most valuable resource. If you need help, the information on page 105 will point you in the right direction.

There's no turning back from here on. If you're prepared to try out this alternative operating system, and are willing to learn new things along the way, turn over the page to begin your journey. You won't look back.

The Pocketbook team
www.pocketbooks.net.au



Installing Linux

CHOOSING A DISTRIBUTION 22

PREPARING TO INSTALL LINUX 23

PARTITIONING YOUR DRIVE 25

MAKING SPACE FOR LINUX 28

BOOTING INTO LINUX 32

INSTALLING RED HAT LINUX 7.0 34

INSTALLING MANDRAKE LINUX 7.2 40



Choosing a distribution Red Hat 8.0 and Mandrake 9.0 are two of the most popular, feature-rich distributions available today. But which one interests you most?

If you have the time and the space, you might as well try them both out, as you'll get a good idea of how distributions can differ from one another. The following may help you decide which you prefer to use in the long term.

RED HAT 8.0

Red Hat's distribution is aimed at stability and performance for usage in business. This results in a more conservative release schedule than Mandrake — Red Hat generally doesn't like to jump the gun on some of the features that Mandrake uses until they have been fully tested. That said, Red Hat's distribution has a number of excellent and unique features, such as its wonderfully easy Internet Connection tool. Beyond the desktop, Red Hat also has the largest support presence in Australia, and the largest body of supported commercial applications. Red Hat was the first distribution to try to make Linux easy for the masses, and the Linux Pocketbook has carried Red Hat since the start.

RedHat 8.0 includes: Kernel: 2.4.18 Xfree86: 4.2.0 Gnome: 2.0.1 KDE: 3.0.3

MANDRAKE 9.0

Mandrake began life as a slightly modified version of Red Hat, but now bears little resemblance. Mandrake's primary goal is to make Linux as easy to use as possible. This is achieved by staying on top of the latest technologies and implementing them as fast as it can. Mandrake was the first Linux distribution to allow partition resizing in the installer, the first to optimise the entire distribution for Pentium processors (providing visibly faster performance) and one of the first to offer a journaling filesystem (Resierfs) as standard. However Mandrake still maintains full compatibility with Red Hat, which means you can install Red Hat and Mandrake packages on either system.

The ease of use and performance Mandrake strives for is the reason it was selected as the second distribution for The 2003 edition of Linux Pocketbook.

Mandrake 9.0 includes: Kernel: 2.4.19 Xfree86: 4.2.1, 3.3.6 Gnome: 2.0.1 KDE: 3.0.2

Preparing to install Linux Before you go about installing your chosen distribution, there are a number of steps that you first need to follow.

Welcome to one of the most exciting aspects of Linux — installation. While operating systems such as Windows splash the screen with images of over-zealous family members in various technology-induced poses, Linux takes a more interactive and informative approach during installation by letting the user know what's going on and giving them full control over how the system is installed and configured.

Many moons ago installing Linux was an adventure that required late nights and tubs of coffee. Now, sadly, those exciting days are over. Graphical installation routines are so slick that Linux is now easier and quicker to install than Windows — quite a feat for an operating system long regarded as a toy reserved for the technically enlightened.

However, while installing Linux is now a simple affair, making sure your machine is in a fit state — especially if you're planning to run both Linux and Windows — requires a little forward planning. The following pages are a guide to help you understand what installing an operating system is all about and to provide

you with the information you need to make the most of your machine.

Many moons ago installing Linux was an adventure that required many late nights and tubs of coffee.

BE PREPARED

Even the most experienced technical gurus of our time make mistakes with their machines. What separates the power users from the common folk is understanding the necessity of prevention. Being prepared maximises the chance of recovering from anything untoward that might occur, and helps you understand what's going on with your beloved machine. So the first step before installing any operating system is to back up any information you deem important. There's an easy test to see if information is important to you. Just ask yourself this question: 'If the contents of my hard drive were to vaporise this instant, what would I miss?'

Yes, it's a hassle when all you really want to do is dive in and play, but while it's extremely rare for hardware and sometimes even



software such as Linux to suddenly trash your information, people aren't so reliable, and mistakes do happen.

So back up.

If you're like most of the world's computer users, you probably just nodded your head at the above and are, at this very moment, sagely ignoring the advice. Don't. You can either learn from your mistakes, or follow this guide and prevent them.

Once you've done the deed and saved valuable information to your media of choice (CD-ROM, Jaz, Zip, another hard drive, even the humble floppy), the next step in being prepared is to read this guide through before you actually start installing, to help familiarise yourself with the procedure. Unless you're a power user, of course, in which case you won't need to read the following pages and won't need to back up — because you've already done it, right?

WHAT YOU NEED

In the past we listed patience and caffeine as necessities for installing Linux. These days all you need to apart from having a linux distro handy is between 1G and 1.5G (or more, if you like) of free hard disk space. In fact, if you can afford the space, you should allocate up to 2G to allow for future growth. To make things easier, we're going to show you how to free up this space if you don't already have it available.

We also used to recommend that you write down important information about your machine to help you answer any queries Linux might have about your hardware. This, too, is no longer a requirement; Linux not only detects hardware efficiently, but can also detect and configure any new hardware you add in the future, automatically.

The only information you might want to keep handy is anything specific to the configuration of your machine — such as an IP address on a network, or the connection details of your ISP. In other words, the type of information that is specific to how you use your machine, which you may have tweaked and expanded over time.

If you're a seasoned Linux user or geek, you can safely bypass the following pages and jump right into installing the distribution of your choice, starting on page 34. Otherwise, feel free to read the following sections, which first cover the nature of drive partitioning and how it relates to installing a new operating system, and then how to specifically make room for Linux on your machine.

Partitioning your drive One of the most important steps when it comes to installing a new operating system is the art of partitioning.

You only need to read this section if you plan to create a dual-boot system with both Linux and Windows. If you are completely replacing Windows or you're installing Linux on a bare machine, you can jump straight to 'Booting into Linux' on page 32.

Linux will happily coexist with other operating systems on your machine, including Windows. In fact, to make dual booting between the OSes easier, Linux will automatically install a boot manager after installation which will let you boot both Linux and Windows. All you need to do before installation is decide where you want to install Linux on your system. To help you decide this it's important to understand how information is allocated on your hard drives.

At the simplest level, your hard drive can be divided up into a number of partitions, which can essentially be thought of as virtual hard disks. One or more of these partitions can reside on your hard drive, and these are used to organise and separate data. For example, a 4G drive might contain two 2G partitions. One 2G partition might contain Windows while the other

might be a storage drive. Although you physically only have one 4G drive in your PC, operating systems such as Windows will see the two partitions as two individual hard drives. Most PCs preinstalled with Windows are likely to have just one partition occupying the entire contents of the hard drive.

The process of partitioning is all about making space to install another operating system, or to set aside a part of your hard disk for storage, that exists independently of areas that contain operating systems. Partitioning isn't hard to do, but there are specialised tools you need to properly manage your partitions.

Linux will happily coexist with other operating systems on your machine, including Windows.

Windows has its own partitioning program called FDISK, however FDISK has a nasty habit of destroying all the data on partitions you modify. Using FDISK is fine if you want to destroy existing partitions to make room for new ones, but we'd rather



you didn't have to do that, so we're going to show you how to resize, and thus keep, existing partitions.

There are many commercial products you can buy to do this, such as PowerQuest's Partition Magic. If you have such a program, we recommend you use it as you follow this guide. If you don't have one, don't fret — Linux comes with its own partition resizing tool.

Before covering the art of partitioning, it helps to understand the different types of partitions and where they are used.

PARTITION DEFINITIONS

PRIMARY PARTITION

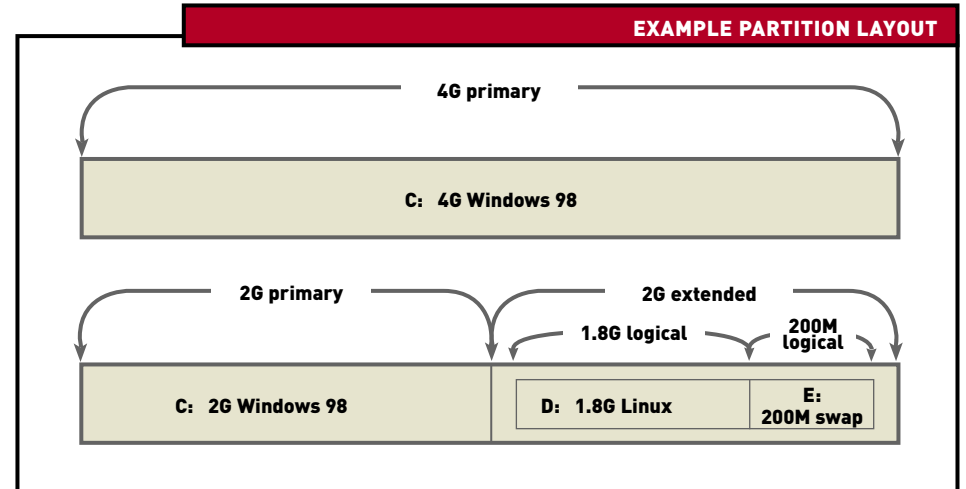
A primary partition will always be the first partition on a hard disk, hence the name. Most machines shipped with operating systems will use just one large primary partition that fills the entire space of the hard drive, and for most users this is acceptable. However, Linux requires its own partition, as well as an extra partition set aside for its swapfile (the file used for virtual memory). If, for example, you currently have Windows installed as your only OS, this will mean breaking down that single partition into three partitions (one for Windows and two for Linux). It's possible to define up to four primary partitions on a hard drive, but this isn't the maximum number of partitions you can create.

EXTENDED PARTITION

An extended partition is basically a workaround for the four primary partition limitation. In itself, it cannot store data. It merely acts as a container in which other partitions, called logical partitions, can be created.

LOGICAL PARTITION

Logical partitions are functionally equivalent to primary partitions. They need a filesystem to be created on them before they can be used (such as FAT32 for Windows 95/98, NTFS for NT/2000 or Ext2 for Linux) and, like primary partitions,



can be used to store and boot operating systems. The only difference between a logical and primary partition is that logical partitions exist only in the extended partition.

ACTIVE PARTITION

The active partition is a 'flag' setting normally assigned to the first primary partition created on your machine. Your PC will always look to the 'active' partition to boot. Only primary partitions can be marked 'active' and only one can be active at any time. If you have only one operating system on your machine, it makes sense that this operating system resides on the first primary partition, as with Windows (you may have noticed that Windows must install its boot files to C:). If you plan on using multiple operating systems, it makes more sense to boot a program that allows you to choose

between operating systems. In this case, the partition where this program resides should be marked active.

To help you make sense of these terms, the example above shows how a single 4G Windows partition can be broken down into a 2G partition for Windows, and a 2G partition for Linux and its swapfile. As you can see, partitions are allocated letters to make it easy to identify them. Note, however, that only DOS and Windows use the drive letter-naming convention. Operating systems such as Linux reference partitions in a different way — something you'll learn more about later. For now you need to choose a partition (or, under Windows, 'drive') with enough free space to install Linux. Unless you happen to have a very large drive or multiple hard drives, this will likely be the C: drive.

THE ALLOCATION OF DATA

It's important to differentiate between a partition and a filesystem. A partition is where information is placed on a drive. As far as the operating system is concerned, a partition is a single, physical hard drive. Your operating system can't use a hard drive until a partition has been defined.

By contrast, a filesystem is how information is stored on a partition. A filesystem determines the structure of data written to the drive and exists within the confines of the partition.



Making space for Linux

Now you know the basics of partitioning it's time to set aside space to create your partitions for Linux.

An average Linux installation will occupy approximately 1.5G, depending on what you choose to install. While a full installation can require more disk space, you don't need to install everything that comes on the distribution CDs. And if you're setting up Linux with a specific purpose in mind (such as a file server or gateway), you can install considerably less. We recommend to set aside 1.5G for Linux and 128M for its swapfile. To allow for future growth it's good to allocate up to 2G, then you won't have to worry about running out space.

Making space for Linux is perhaps one of the most confusing tasks for new users. It is, however, the only hurdle to overcome. Once your drives are partitioned, the rest is smooth sailing. Even if you need to reinstall Linux later, your partitions will already be set up, so you won't have to do it again.

In order to make space for Linux you need to create new partitions out of the free space on existing partitions. This is known as resizing. As in the example of the 4G drive (see previous page), you could reduce the primary 4G partition down to 2G, thereby leaving 2G of space to create partitions for Linux. Before you can resize, however, there are a number of steps you need to follow.

1 STEP ONE

BACK UP YOUR DATA

Not even commercial partitioning software can guarantee something won't go wrong, and although partitioning is a quick and simple process, a simple mistake can destroy all the data on your hard drive. We've already covered the importance of backing up, so if you haven't done it yet, do it now before reading on. From this point onwards you're actually going to be working with your drives.

2 STEP TWO

CREATE FREE SPACE

To resize you need to shrink existing partitions to make room for new ones. You can't do this unless you actually have free space on your existing partitions.

So the next step is to make sure you have enough free space equivalent to the amount of space you want to give to Linux. If you're short on space this will mean deleting programs, tools and games you no longer need or use.

3 STEP THREE

DEFRAGMENT YOUR PARTITIONS

If you know a little about file allocation under Windows you'll know that files can be written anywhere on a drive. Over time, with programs coming and going, data can be scattered all over a partition. This is called fragmentation. Resizing software can't reduce a partition if any data lies near the end of it, so it's necessary to defragment your drive in order to move all the data to the front of the partition. Under Windows simply click Start → Programs → Accessories → System Tools → Disk Defragmenter to defragment all of the partitions you plan to resize.

THE ART OF PARTITIONING

There are a number of different methods you can use to resize partitions to make space for new ones:

- ✦ Use a commercial tool such as Partition Magic;
- ✦ Use the resizing feature of Mandrake Linux;
- ✦ Use the open source FIPS program, included with both Red Hat and Mandrake Linux.

If you've decided to install Mandrake Linux you can resize your partitions during installation, so jump ahead to 'Booting into Linux' on page 32 to start installing Mandrake Linux.

If you've decided to install Red Hat we recommend you use a tool such as Partition Magic over FIPS, simply because Partition Magic is easier to use.

Alternatively, you could start a Mandrake installation, make space for Linux using its tools, and then restart with a Red Hat installation. For the sake of tradition, or in case you're really keen to use FIPS, we'll show you how below.

USING FIPS

FIPS, the First Interactive non-destructive Partition Splitting program, isn't as fancy or as feature rich as a tool like Partition Magic, it doesn't even use a GUI. However, like the rest of Linux, FIPS is free.

Technically, FIPS isn't so much a resizing tool but rather a partition splitter. It will take a single partition and split it into two, creating a new, empty partition out of the free space at the end of the old partition. The new partition can then be used to install a new operating system, or in the case of Linux, it can be deleted and two more created in its place (one for Linux and one for the Linux swapfile). Don't worry about creating partitions yet; that's easily done during the installation of both Red Hat and Mandrake Linux.



Before you start, write down the size of the partition (or drive) you intend to split as well as the size of the new partition you want to create. If you're following our earlier recommendations, this will be around 1.5G.

1 STEP ONE

CREATE A BOOT DISK

The first step is to grab a spare floppy (you haven't quite abandoned floppies yet, have you?) and make it bootable by opening up a DOS prompt under Windows and typing:

```
FORMAT A: /S
```

2 STEP TWO

COPY FIPS TO YOUR BOOT DISK

Next, copy the following files from the Red Hat or Mandrake Linux installation CD:

```
FIPS.EXE
RESTORRB.EXE
```

You'll find both these programs in the \DOSUTILS\FIPS20 directory on the CD. The second file can be used to recover your old partition information which you will be prompted to save by FIPS. You might also want to read some of the FIPS documentation from this directory, particularly ERRORS.TXT, FIPS.DOC and FIPS.FAQ. There's no harm in familiarising yourself with the program that is about to handle your precious data.

3 STEP THREE

STARTING FIPS

If you're ready to proceed, boot your floppy disk. When you arrive at the A: prompt, start FIPS by typing `fips`. FIPS should automatically detect your drives and, if you have more than one, ask which one you want to partition. FIPS will then display lots of information about your hard disk. It doesn't matter if you don't know what it all means; just look at the 'MB' column to identify your partitions by their size in megabytes.

4 STEP FOUR

CREATING A NEW PARTITION

If FIPS is able to read your file allocation table without any problems, it will then ask you which cylinder the new partition should start on. As you may be aware, a hard drive's capacity is measured in cylinders and drive heads, and if you're a particularly technical user, you probably know how cylinders translate to disk space. Fortunately, you don't need to know this; simply use the left and right arrow keys to adjust the number of cylinders for your new partition and you'll see the estimated size in megabytes change accordingly.

Although you will eventually need to use two partitions (or more, if you wish) for Linux, this will be automatically handled by the installation procedure. For now, all you need to do is create a partition equal to the total amount of space

THE LAW OF CYLINDERS — A NOTE ABOUT LARGE DRIVES

Many operating systems, including Linux, have problems booting from partitions that exist above the 1,024 cylinder boundary. This boundary, an inheritance of ancient PC architecture, often causes problems for users who install multiple operating systems. As a result, it's important that you install Linux below this boundary or, at the very least, install the Linux boot files (/boot) below this boundary.

If this sounds a little techy, the shorthand is to try and ensure the Linux boot files reside on a partition within the first 4G of a hard drive. Alternatively, Mandrake Linux will enable 'lba32' boot support automatically during installation, allowing you to install Linux above the 1,024 cylinder boundary.

you want to allocate to Linux. Try to allocate more rather than less, so you have room to expand in the future.

5 STEP FIVE

SAVING YOUR CHANGES

Once you're happy with the values, press Enter. Follow the instructions (re-edit the table if necessary, or press 'C' to continue), then write the partition table to disk.

Note that FIPS isn't a particularly user-friendly program. Take your time; read the instructions onscreen carefully and make a backup of your partition table when prompted. If anything goes wrong, you can restore your previous partition layout. Also note that changes aren't written to disk until you say so. Don't worry if you make a mistake adjusting values; just correct them or start again. Only write the changes when you're absolutely certain of the partition layout.

6 STEP SIX

YOUR NEW PARTITION

Now reboot to Windows and make sure everything works well. You may have noticed that certain drive letters after C:, such as the letter of your CD-ROM, may have changed. This is perfectly normal and will be rectified when your Linux partitions are created later. Windows can't see Linux partitions, and as such, they won't be mapped to a drive letter.

You have now successfully created a new primary partition using FIPS out of the space you made available on the old partition. This is about as hard as it gets. You can relax now, grab a cup of coffee and get ready for the most exciting step — installing Linux!



Booting into Linux

When it comes to starting a Linux installation you have a number of choices available. Booting from the installation CDs is the best method, but there are other options.

The first and ideal method to start installing Linux is booting the distribution CD, as this allows you to install Linux with a minimum of fuss. If this isn't an option, be sure to try the autorun features listed below before you fall back on that old faithful medium, the floppy disk.

BOOTABLE CD

Both Red Hat 8.0 and Mandrake 9.0 are bootable CDs, and booting them is the easiest method for installing Linux.

If you have a machine capable of booting CD-ROMs you may still need to tell your PC that you want to do so. Today, most BIOSes have an option to choose the order in which devices will be booted. For most machines, this usually involves checking A: (or the floppy drive), followed by the hard drive. If your BIOS supports booting from a CD-ROM, you'll be able to add your CD-ROM to the queue and place it first. If you're not sure how to enter your BIOS or don't want to fiddle with its options, or your machine isn't able to start from bootable CD-ROMs, you'll need to follow another procedure, listed below.

AUTORUN FROM CD

Both Red Hat 8.0 and Mandrake 9.0 are capable of starting Linux and then launching the installation procedure directly from the CD.

To do this you first need to make a DOS boot disk that has support for your CD-ROM drive. If you don't already have one, the easiest way to make one is to use Windows 98's Startup Disk feature. To do this Click Start → Settings → Control Panel → Add/Remove Programs → Startup Disk and then create one. The boot disk Windows 98 creates uses a generic IDE CD-ROM driver that supports most IDE CD-ROMs, and provides support for a wide range of SCSI CD-ROMs.

When you're ready to install, simply boot the floppy disk, switch to your CD-ROM drive (likely to be mapped to D: or higher), then type:

```
Red Hat 8.0
CD DOSUTILS
AUTORUN
```

```
Mandrake 9.0
CD DOSUTILS/AUTOBOOT
AUTOBOOT
```

Starting the installation procedure this way is functionally equivalent to booting the CD.

INSTALLATION FLOPPIES

Red Hat 8.0 and Mandrake 9.0 also support starting the installation from a boot floppy, although this should be used as a last resort. You will need one blank floppy disk. To save you time and trouble later on, make sure the disk is at least relatively new and reliable, and not the dusty disk graciously supporting the coffee mug on your desk.

To create the boot floppy, boot to DOS or simply open a DOS command line under Windows, change to your CD-ROM drive and type:

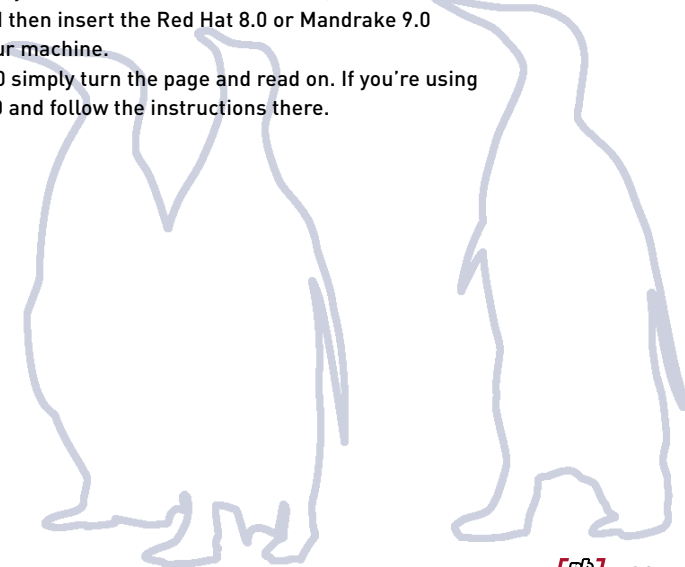
```
Red Hat 8.0
CD DOSUTILS
RAWRITE /IMAGES/BOOT.IMG
```

```
Mandrake 9.0
CD DOSUTILS
RAWRITE /IMAGES/CDROM.IMG
```

STARTING THE INSTALLATION

When you're ready to start the installation make sure you've followed the instructions detailed above for your chosen installation method, insert a boot floppy if you're using one, and then insert the Red Hat 8.0 or Mandrake 9.0 installation CD and reboot your machine.

If you're using Red Hat 8.0 simply turn the page and read on. If you're using Mandrake 9.0, turn to page 40 and follow the instructions there.





Installing Red Hat Linux 8.0 Red Hat has been around almost since the beginning, and is one of the most popular distributions in the world today.

After you've booted Red Hat Linux using your chosen method you should find yourself at the boot prompt. Press Enter, and the install program will load.

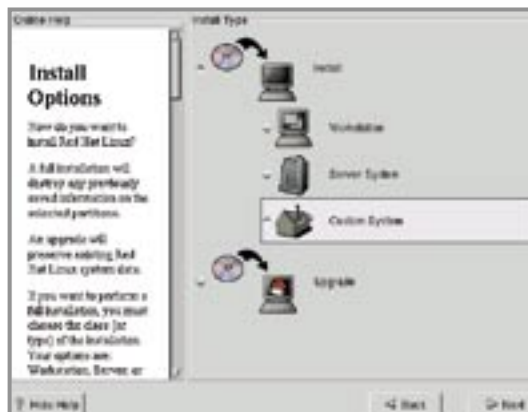
Because the installation process is now so easy, we'll cover only those aspects of the routine that need more explanation.

The process begins with a few questions about your language, keyboard and mouse. The default values for the language and keyboard should be fine, and the type of mouse is usually detected correctly. However, you should make sure the correct number of buttons is selected (if you have a two-button mouse, ensure that 'Emulate 3 buttons' is turned on).

INSTALLATION TYPE

When prompted for the installation type, your choices are Workstation, Server or Custom. If you already have an older version of Red Hat on your system, you can also choose to Upgrade it. The Workstation and Server installs can automatically create the partitions where Red Hat Linux will be installed, but the defaults may not be appropriate. Choose Custom, as this offers more flexibility.

Choosing an installation type.



PARTITIONING

After choosing the Custom install, you can use Disk Druid or fdisk to do your partitioning. Unless you have very specific requirements for your partition layout, choose Disk Druid; it's far easier than the antiquated fdisk.

Disk Druid will search your system and display a list of existing partitions (in the top window) and hard drives (in the bottom window). You need to create two partitions for Linux to use: one for the root filesystem and another for the swap partition.

If you've followed the previous section and used fips to free up space, you should find a nice, large partition ready to be deleted.

THE ROOT FILESYSTEM

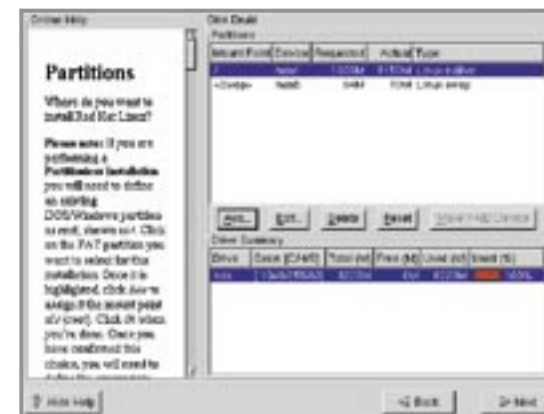
All of the directories in Linux branch off one main root directory (you'll learn more about this later in the book), so you need to create this partition first.

To add a new partition, select Add. Type a single forward slash (/) as the mount point. Set the partition type to Linux Native, and the size to whatever is appropriate for your system. If you want to use the remaining space for Linux, simply type in a small number (such as 200M) for the size and select the 'Use remaining space' checkbox.

Now click OK, and a new root partition should appear.

THE SWAP PARTITION

Choose Add again, and this time select the partition type 'Linux Swap'. There's no need to enter a mount point. Sizewise 128M is plenty for most installations. If you activated



'Use remaining space' for the root partition, make sure it is turned off for the swap partition, and then click OK.

Disk Druid in action.

Because space is allocated by cylinders, which are generally larger than 1M, the value in the 'Actual' column will rarely match the exact value you requested. If you want to change the amount of space allocated to a partition, just select the partition and click Edit. Any partitions that you create can be changed at any time — until you click Next to continue with the installation.

WHAT IS A MOUNT POINT?

A mount point is the location where Linux adds (or mounts) a filesystem to the directory structure. For example, the root filesystem you create will be mounted as '/' (the root directory).


Any other filesystems — such as floppies, CD-ROMs or other hard disk partitions — that you want to use will be mounted in different directories on different mount points. This is different from DOS and Windows, where each device effectively has its own root directory and is distinguished by letters.



LINUX DRIVE NAMING

When partitioning, you'll notice that Linux doesn't label its drives and partitions with letters (eg. C:, D:, and so on) like Windows does. Instead, drives and partitions are just another form of 'device' to Linux. As a result, hard drives and partitions are named as follows:

IDE drives always begin with '/dev/hd', SCSI drives with '/dev/sd'. The drives are labelled 'a' for the first drive, 'b' for the second, and so on. The partitions on each drive are numbered starting from '1', but you may find that some numbers are skipped — the first four are reserved for primary partitions, and '5' onwards are for logical partitions (see page 25 for more information on partitions). So, if you have a single IDE drive partitioned into Windows, Linux native and Linux swap partitions, you might see something like this:



/dev/hda1	Windows FAT partition
/dev/hda2	Linux Native
/dev/hda5	Linux Swap

When you do click Next, you will be asked if you want to format the partitions you created — this is a necessity when doing a first-time install, so leave the box checked. You can also choose to check for bad blocks, a step that ensures the disk is in good condition. Note that this option can take a long time to complete, and isn't really necessary with newer machines.

LINUX LOADER (LILO)

Next, you'll need to specify where the Linux boot manager, LILO, is to be installed. LILO allows you to boot other operating systems — including DOS and Windows — and unless you want to use a boot floppy every time you start Linux, you'll need to install it (even if you use a commercial boot manager such as System Commander). There are two places where you can install LILO:

- ☑ in your Master Boot Record (MBR)
- ☑ in the first sector of your boot partition

Generally, install LILO in the Master Boot Record if you don't have a boot manager. Otherwise, install it in your boot partition and follow the boot manager's

instructions for adding a new entry to your boot manager software.

The LILO configuration program will look for any other bootable partitions, such as an existing Windows installation, and automatically add them to the LILO boot menu. You can customise this menu, or simply accept the defaults. We recommend leaving the 'Create Boot Disk' option selected — if there are any LILO boot problems, you can boot this disk and be able to access your Linux system.

NETWORKING

To network your computer, fill in your IP address and other necessary details. If you don't have a gateway or a DNS (domain name system) server, leave those fields blank. If you have a DHCP (Dynamic Host Configuration Protocol) server on the network, Linux will automatically find out all it needs to know from the DHCP server.

If you don't have a network card, you can still set your machine's hostname now to give your machine a more personable name than 'local host'.

TIME ZONE

Selecting the time zone is one of the more impressive aspects of a Red Hat Linux installation. By clicking near your location on the world map, time zone information will automatically be configured. If you find it difficult to point precisely, you can zoom in on the Pacific Rim area.

The option for 'System clock uses



UTC' should be left unchecked if your machine will be dual-booting with Windows.

Selecting the time zone.

PASSWORD

As Linux is natively a multiuser system, you need to choose a password for the root account (otherwise known as the system administrator account). You can easily change the password later, so don't worry too much if you can't think of a good password immediately. But don't, under any circumstances, forget it! It is possible to recover from a lost root password, but it's not the type of thing you'll be keen to learn about at this stage.

You can also create a regular user account while you're here. You should do this now, as the root account should only be used for administrative tasks. The account you create here should be your everyday use account.



AUTHENTICATION CONFIGURATION

These protocols are used to ease administration of user details on large networks. Unless you know that you are on a network that uses NIS, LDAP or Kerberos, just leave these options on the defaults and click Next.

CHOOSING PACKAGES

Welcome to the best part of installing: choosing packages! What's a package? A package is a method of storing a bundle of related files together in a way that is easy to install and remove.

Package selection.

As a file format, it is similar to Zip, but unlike Zip, Linux packages contain all sorts of information that allow you to catalogue which packages are installed on a system, so you can uninstall them at the click of a button.

There are over 850 packages in Red Hat Linux, and to make it easier to install the ones you want, Red Hat has grouped them into 'package groups'. The groups we recommend that you install are:

- You may also want to install the following package groups, depending on your
- X Window System
- Gnome
- KDE
- Mail/WWW/News Tools
- DOS/Windows Connectivity
- Graphics Manipulation
- Multimedia Support
- Networked Workstation
- Dialup Workstation
- Kernel Development
- Utilities

needs and interests:

- Printer Support
- Games
- Laptop
- SMB (Samba) Server
- IPX/NetWare Connectivity
- Anonymous FTP Server
- Web Server
- SQL Server
- Network Management Workstation
- Authoring/Publishing

Most of the package groups are self-explanatory, except for perhaps 'SMB [Samba] Server'. If you have a network with Windows machines on it, then

you will probably want to install the Samba server as it lets you share files and printers from your Linux machine using the standard Microsoft Network protocols. (For information on how to use Samba see page 138.)

At this point you have the option to 'Select Individual Packages'. We recommend you choose this and then, on the following screen, spend some time just exploring all the packages that come with Red Hat Linux. Selecting a package will display a description of its purpose. If you decide you want to install it, just double-click on it. Packages of interest include:

Applications → Editors → Abiword
An easy-to-use word processor

Applications → System → Procinfo
Displays summary information about your system, including average CPU, memory and interrupt statistics

User Interface → X → 100dpi-Fonts
Higher resolution fonts for high-resolution monitors

THE X WINDOW SYSTEM

The last step before copying packages is to configure the X Window System (the GUI that powers the Linux desktop). Generally, the installer can automatically detect your video card and monitor settings, but in some cases it might not be able to. If this happens look for the 'Generic' monitor model and choose a monitor definition that matches the capabilities of your own. Similarly, if your video card isn't



detected, select 'Generic SVGA'.

Make sure 'Use Graphical Login' is checked and that Gnome (or, if you prefer, KDE) is your Default Desktop (you'll learn more about these environments soon).

Be sure to click the 'Customise' button before you proceed so you can tailor the video modes to your liking. Unless you plan to switch resolutions, just select the highest resolution for the highest colour depth your video card can display. Select 'Test this configuration' to make sure they work.

Once you're happy with the X configuration, you're ready to start installing Linux. Depending on what you've selected to install, you may be prompted to insert the second installation CD. Installing can take anywhere from a few minutes to half an hour, depending on the speed of your CD-ROM and other components of your system. Now would be a good time to kick back with your favourite caffeine-infused beverage and watch the fun!

Choosing individual packages.



Installing Mandrake Linux 9.0

Mandrake is a relative newcomer compared to Red Hat, but already has a loyal following due to its emphasis on ease of use and performance.

After you've booted Mandrake Linux using your chosen method, you should find yourself at the Mandrake SYSLINUX boot prompt. Press Enter and the graphical install program will be loaded. In some cases (for example, if your system doesn't have sufficient memory), you may want to load the text-based install program instead. To do this, type text then press Enter.

The installation procedure for both graphical and text modes is similar, so you should have no problem following the instructions below regardless of the method you choose. Because the installation process is now so easy for Linux, we'll cover only those aspects of the routine that need explanation.

Throughout the installation, a list of the steps that need to be performed will be displayed on the left-hand side of the screen. The stars next to the steps are colour coded to indicate which steps have been performed already, and clicking on a star will take you directly to that installation step. You'll also notice four thin colour bars at the bottom of the screen; clicking on these will change the colours for the installation routine. Try them out — it makes installing Linux a little more aesthetically pleasing.

Choosing an installation class.



INSTALLATION CLASS

After choosing your language, you'll be prompted for an installation class. Which class you select will affect how many questions the install program will ask you later on and how much choice you get in customising your system for the installation.

For the purposes of this guide we'll walk you through a 'Customised' installation, as this offers more flexibility than the 'Recommended' installation.

SYSTEM USAGE

The Mandrake installer will ask whether your system will be used as a Workstation, Server or for Development. This selection is used to preselect specific defaults for the installation. For the purposes of this guide select Workstation, as this contains all the applications and configuration for a general-purpose Linux installation.

MISCELLANEOUS QUESTIONS

The installer has a few miscellaneous questions which don't fit anywhere else in the installation. These are:

- ☞ Hard drive optimisations
- ☞ Security level
- ☞ Precise RAM size
- ☞ Removable media automounting
- ☞ 'Num Lock' at startup

Turning on hard drive optimisations can improve the performance of your disks by enabling faster transfer modes. Unfortunately, there are some controllers and drives that claim to support these modes, but cause data corruption instead. If you have a newer system, these optimisations should work just fine and will improve your disk performance.

There is often a tradeoff between security and ease of use, so Mandrake allows you to choose what security level you want. Low is fine for a standard home use system, but Medium is a good compromise if you want a more secure machine. If you were to install in Expert mode you would have the option

of setting security as low as 'Welcome to crackers' and as high as 'Paranoid'. Choice is a wonderful thing!

Setting the precise RAM size is no longer necessary, but is useful for those rare occasions on very old machines which don't properly report the amount of RAM available. Don't change this setting unless you're certain the amount of RAM Mandrake has detected is incorrect.

The removable media automounting ('supermount') feature allows you to use CD-ROMs and floppy disks without having to mount and unmount them. Make sure this option is selected, as it makes using removable media much easier under Linux.

PARTITIONING

When it comes to partitioning you'll be presented with a number of different choices, depending on your partition layout. Select DiskDrake, as this is by far the easiest option and gives you full control of how and where Linux will be installed.

DiskDrake is a rather nifty program that not only allows you to create and delete partitions, but also resize them. If you've already made space for Linux you can simply click on the partitions you've made available and set the root (/) and swap mount points. Otherwise, this guide assumes you have a single, large Windows partition ready to be resized after following the instructions in 'Making space for Linux' (such as defragmenting).



Start by selecting the Windows partition and clicking on Resize. A dialog will appear asking what new size you would like the partition to be. Type in a value, or use the slider bar, to reduce the Windows partition by the amount of space you'd like to allocate to Linux.

Once completed you should have a large blank area visible in the partition diagram. You can try out the 'Auto allocate' function to let Mandrake decide how to best break up the free space for use by Linux, but we recommend creating two partitions, one for

DiskDrake in all its glory.

root and one for swap, for the sake of simplicity. Click on the blank space and select 'Create'. Type in (or use the slider) the amount of space you want to allocate to the root partition. If you're following our recommendations this will be between 1.5G and 2G. Remember to leave enough space for the swapfile partition — up to 128M. If you can afford 2G for the root partition, go all the way as this will allow you to install more features as well as leave room to grow in the future. Select 'Linux native' as the partition type and then select '/' as the mount point.

Next, click on the remaining space and repeat the procedure, this time creating a swap partition. Note that you won't need to set a mount point for the swap partition.

If you can afford 2G for the root partition, go all the way as this will allow you to install more features.

You may also notice in the partition creation dialog the option of setting the filesystem type to 'ReiserFS' and 'Linux RAID'. ReiserFS is a new journaling filesystem — still under development

— that offers superior performance and filesystem integrity over the standard Linux Ext2 filesystem. Try it if you want, or stick with Ext2 if you're not sure. The Linux RAID option allows you to create software RAID arrays to increase speed, data protection, or both. To make use of it you will need at least two physical hard drives. (Linux RAID and ReiserFS will be covered in more depth in the 2003 edition of Advanced Linux Pocketbook.)

After partitioning your disk, any new partitions that you have created will be formatted; depending on the size of your disk, this should only take a few moments. After this, Mandrake should present you with a prompt asking if you have the second Extensions CD; select Yes.

CHOOSING PACKAGES

The next step is to choose the packages you want to install. A package is a way of storing a number of related files so that they can be easily installed and removed together. Almost all of Mandrake has been divided up into packages and, generally, each piece of software included with Mandrake is placed into its own package for ease of management.

The packages have been further grouped into 'package groups'. You'll now be presented with the option of choosing your package groups. We recommend you install the following packages:

- KDE
- GNOME
- Internet tools
- Communication facilities
- Multimedia
- Graphics Manipulation
- Documentation
- Development C/C++
- Console Tools
- Utilities
- Miscellaneous

You may also want to install these package groups, depending on your interests:

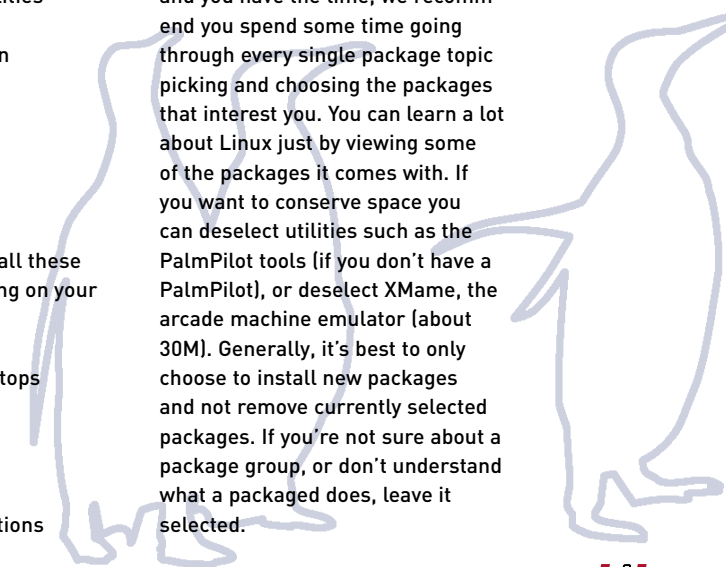
- Other Graphical Desktops
- Office
- Games
- Databases
- Development Other
- Development Applications

Before continuing, make sure that 'Individual package selection' is highlighted.

The next screen will present you with the option of prioritising the number of packages installed if you don't have enough space to install all the selected package groups. If you don't have enough space we recommend going back a step (click on the appropriate star on the left-hand side) and unselecting a package group.

If you're curious and you have the time, we recommend you spend some time going through every package topic.

The following screen will display a list of all the packages contained in Mandrake 9.0, organised in a tree hierarchy. If you're curious enough and you have the time, we recommend you spend some time going through every single package topic picking and choosing the packages that interest you. You can learn a lot about Linux just by viewing some of the packages it comes with. If you want to conserve space you can deselect utilities such as the PalmPilot tools (if you don't have a PalmPilot), or deselect XMame, the arcade machine emulator (about 30M). Generally, it's best to only choose to install new packages and not remove currently selected packages. If you're not sure about a package group, or don't understand what a packaged does, leave it selected.





Some packages you may want to add include:

Games → Arcade → Chromium

A nifty OpenGL scrolling shoot-em-up game. Not just a bit of fun but also handy to test the OpenGL support for your video card. If it runs slow, you know support isn't installed and you can seek out the appropriate OpenGL drivers for your card online.

Development → Kernel → kernel-source

This is the source to the Linux kernel. We recommend you install this if you're interested in exploring the kernel or maximising the performance of your system. Compiling your own kernel is covered on page 117.

Graphical Desktop → Other → GTK Themes

More GTK themes for Gnome than you can poke a theme manager at.

System → Configuration → Fonts → X11bitmap → XFree86-100dpi

High-resolution fonts for desktops running in high resolutions (1280x1024 and above).

System → Base → Servers → anontftp and telnet-server

Two packages providing anonymous FTP logins and a Telnet server. Not installed by default for the Workstation class, but handy if you plan to use your machine remotely.

Mandrake installing.

When you've finished selecting packages, you can click Install and then sit back with your favourite herbal pick-me-up drink and watch TV while Mandrake copies packages to your system.



CONFIGURING NETWORKING

After copying packages Mandrake will prompt you to configure networking. Mandrake supports the configuration of both dialup and Ethernet networking at installation time.

To configure dialup networking, you'll need to specify what serial port your modem is connected to, and then the connection name (the ISP's name will suffice), phone number, login name,

password and any DNS settings (though DNS is likely to be automatically configured for you when you connect). You can also set the authentication type, but the default 'PAP' is appropriate for most ISPs.

Installing local networking requires setting a static IP or selecting to use DHCP if you're on a DHCP network. Whichever you use, be sure to set your hostname to give your machine some personality.

TIME ZONE

Time zone is self explanatory. However, don't answer 'Yes' to the GMT question if you're installing a dual-boot system.

PRINTER CONFIGURATION

If you have a printer, either locally attached or available on a networked Linux/Unix or Windows machine, you can configure it during the installation.

Setting up a locally attached printer requires entering the name of the printer ('lp' is the default printer, so you may as well leave it at that), what port it's attached to (usually '/dev/lp0' for the first parallel printer port), and selecting the printer type.

Remote (networked) printers require the same information, as well as the hostname of the printer server and the remote queue name. For Windows-connected printers, you'll also need a share name, username and password.

If it's all too much trouble (after all, printers are boring), you can



leave this until later and configure your printer in Linux.

Configuring networking.

ROOT PASSWORD

As Linux is natively a multiuser system, you must choose a password for the root account (otherwise known as the system administrator account). You can easily change the password later, so don't worry too much if you can't think of a good password immediately. But don't, under any circumstances, forget it! It is possible to recover from a lost root password, but it's not the type of thing you'll be keen to learn about at this stage.

REGULAR USERS

You'll need an everyday account for your own use. Create one here and pick a user icon that takes your fancy. (Yes, penguins really do use Linux.)

Even experienced Linux administrators will do most of their work as a regular user, and only switch to using the root user when it's required.

Mandrake will ask for a ‘real name’ (that is, your real life name) and a ‘user name’, which is how you’ll identify yourself to the system when you login. Click on Add and then Done, unless you want to add some more users to the system.

CREATING A BOOT DISK

A boot disk is one of those things that every installation guide ever written always suggests the reader should create. At the same time, it’s the one single step of the installation procedure that hardly anyone ever performs.

That said, we recommend you create a boot disk. Really. You never know when you might need it.

BOOTLOADER OPTIONS

The default bootloader options are usually fine, but if you already have a boot manager such as System Commander, set the boot device to be the Linux partition. Mandrake will automatically detect and add to the boot menu your Windows partition as well as set up a number of different Linux boots — such as the fail-safe boot, which is handy if you ever need to boot a bare system to fix problems.

Although it doesn’t display it, Mandrake will use the GRUB (GRand Unified Boot Loader) boot manager in lieu of the traditional LILO (Linux LOader). If you perform an expert install you will have the option of choosing between the two. Both are installed regardless, in case you decide at a later date to switch from one to the other.

X CONFIGURATION

Selecting the autologin.

The final step of the Mandrake installation is to configure the X Window System, the engine behind the Linux GUI. Your video card and monitor should be automati-

cally detected, but if not, just select an appropriate card and monitor from the list. You’ll then be prompted for the default resolution and colour depth.

After configuring X you will be prompted by the Default user dialog box. If you want Linux to automatically login with your user account at bootup, select the account name and desired graphical environment here. Leave this blank if you want Linux to boot up into the usual login screen.

After this you will be prompted to remove the CD and then Mandrake will reboot your machine.





Exploring Linux

UP AND RUNNING 48

UNDERSTANDING THE DESKTOP 51

EXPLORING GNOME 53

CUSTOMISING GNOME 57

EXPLORING KDE 60

CUSTOMISING KDE 64



Up and running

You've got an all new supercharged Linux installation, a warm beverage steaming in your hand, and a curiosity to rival the Cheshire Cat. Now what?

The first screen you'll be presented with after rebooting will be the LILO (or GRUB) boot manager. Simply press enter to boot Linux.

Linux will load and display lots of geeky information as it detects your hardware and starts system services. If you're a seasoned PC user, you'll recognise much of the information, though it won't stay on screen for long.

When Linux finishes booting you'll be presented with a login screen (technically called a display manager). Type in the name and password of the user account you created during installation (remember they are case sensitive) and select the desktop environment you'd like to use.

THE ROOT ACCOUNT

Linux is a multiuser system. Apart from the user account you created during installation for everyday use, every Linux system comes with a system administrator account, otherwise known as the root account. The root account (also known as the super-user) is the most privileged of accounts; it has full access to the system and is capable of performing any task. In contrast, regular users

have a small set of abilities by default — to run necessary programs, and save documents and files to their own home directory. To add and remove software, modify system settings, or do anything else that could potentially damage the system, you need to use the root account.

While it's tempting to use the root account for everyday use, don't. The unlimited access it provides means a simple mistake could damage your system, or destroy data, even across other partitions such as your Windows drive. Even Linux masters don't use the root account for everyday use.

SHUTTING DOWN

Now that you've logged in it's important we show you how to log out and shut down before we do anything else. Linux, as with most modern operating systems, needs to be shut down before powering off. Never just turn the power off, always shut down first, or bad and nasty things may happen. There are a number of different ways to shut down:

- ☛ Simply click the Gnome or KDE menu button (we'll introduce you to these two graphical environ-

ments soon) and select Log Off. You will have the option to log out, halt the system (shut down and turn the computer off), or reboot (shut down and reboot)

- ☛ Should you ever find yourself needing to shut down from the console (more on this later), the command to use is shutdown. The syntax for this command is:

```
shutdown [options] [when to shut down]
```

The command `shutdown -h` now will shut down and halt the machine, while `shutdown -r` now will shut down and reboot the machine.

- ☛ If you find this too taxing on the fingers, there is another way. While it's important to know the shutdown command, the following also has the same effect: `reboot`
- ☛ Lastly — but you really shouldn't shut down this way — the three finger salute of yore will also work, but only from the console. Don't get into the habit of shutting down this way, or you'll be laughed at by real users. Stick with the logoff buttons in Gnome and KDE or use the shutdown commands — your Linux system will love you for it.

DIFFERENTIATING BETWEEN WINDOWS AND LINUX

If you've grown up in a Windows world then you'll no doubt be familiar with certain conventions relating to how a computer operates. However, these conventions belong only to Windows — each operating system has its own way of performing certain tasks. Below is a summary of some of the differences you will find between Windows and Linux.

SYSTEM CONFIGURATION

Configuration in Windows is done via the Control Panel, administrative applications or the Microsoft Management Console. These utilities are interfaces which edit a centrally stored database called the registry. The registry is not human readable, but can be manipulated directly if desired through programs such as `regedit`. Changing even the simplest of settings often requires a reboot so the registry can be re-initialised.

Linux configuration is performed via any number of command line, graphical or Web-based applications. Such programs manage human-readable text files centrally located in a special configuration directory (which you'll learn about later). As they're text files they can be manipulated directly if desired. Very rarely will you ever need to reboot after reconfiguring Linux.



GRAPHICAL OPERATION

Windows uses a graphical interface called EXPLORER.EXE, which is required to run Windows. Linux does not require a graphical interface to operate, but does offer a number of different graphical environments to work with if you so choose. Even while using a graphical environment, you still have full access to the underlying operating system. Low memory systems or server machines can save valuable resources by not running a graphical interface.

COMMAND LINE

Windows has a very limited command line interface called COMMAND.COM. Linux provides a number of command line interfaces (called shells), with the power to perform almost any task on the system. However, you can manage and completely configure a Linux system without using a graphical interface, and even do so remotely over a network or the Internet.

FILESYSTEM LAYOUT

Windows installs program files in a mixture of the program's own directory and the Windows system directories. As many Windows users know, this causes system instability and compatibility problems on systems where many applications have been installed. Linux installs files based on their type (documentation, configuration files, and so on) and their role in the system (whether they are part of the core system or a simple end-user application). A user's files are kept in their home directory.

SOFTWARE INSTALLATION

Software under Windows usually comes as an executable installer, and is removed via the Add/Remove Programs application in the Control Panel. Software under Linux usually comes as a standardised 'package' that is installed and uninstalled via a program called a package manager. The package manager also keeps track of the files a package depends on, and ensures that installing and uninstalling applications won't break other programs — a welcome change!

A QUICK LOOK AT LINUX

By now you're probably keen to sink your teeth into Linux to see what all the fanfare is about. Discovering Linux's capabilities is something you're likely never to stop doing as you continue to learn more about it on the Web and through books, magazines, and so on.

The following pages will help you explore your new system further, so fill up that coffee cup, relax, and be prepared to journey into an operating system that puts you in control.

Understanding the desktop The Linux graphical user interface is somewhat more complex than its Windows counterpart, but also considerably more flexible.

You've probably heard a lot of different terms used to describe different aspects of the Linux interface — X, XFree86, Gnome, KDE, window managers, and more. Why are there so many different aspects to the Linux GUI. Essentially, the graphical environment is broken down into three separate components, and there are some brilliant technical advantages in separating the GUI like this.

THE X FACTOR

During the setup process, you may recall when your video card was configured that the installation program referred to something mysteriously known as X. The X Window System, as it's more formally known, is the powerhouse behind graphics under Linux. It provides access to the video card in your system and supplies an underlying graphics layer from which other programs, called window managers can load. As we're about to explore, window managers provide the look and feel of your Linux desktop.

Why separate these layers? Technically, X isn't just a graphical subsystem, but rather a full network-aware client/server graphical protocol. That is, in addition to displaying graphics on your screen, the same X server can generate graphics on any other screen across a network, even at the same time. This capability has been used for over a decade in government, research and educational institutions to allow a centralised server to provide powerful operating environments for its users. The concept has only recently appeared in the Windows world with the use of applications such as Citrix Winframe — but Windows isn't inherently designed to support this concept and doesn't do it terribly well.

Since X is an open standard, many different versions have appeared offering different functionality. The most popular which provides X on Linux is called XFree86.

WINDOW MANAGERS

Window managers handle the placement of windows on screen; how and when they will be raised and lowered, maximised and minimised; and the decorations around the edge of each window. Different window managers provide different



functionality, have different looks and use different amounts of resources. Some window managers such as Enlightenment offer an amazingly high level of configurability in how the desktop looks, but also take up lots of RAM. A small, thin window manager such as IceWM uses far less resources, but doesn't offer the degree of functionality that Enlightenment does.

DESKTOP ENVIRONMENTS

Desktop environments provide drag and drop, object models


Lightweight window managers such as Window Maker can run without a desktop environment.

(eg. cut and paste), session management (eg. prompts to save documents when logging out), icons, buttons and other common GUI features. The most common desktop environments on Linux are Gnome and KDE. Desktop environments they can operate with a variety of window managers. For example, although Sawfish is the standard window manager for Gnome, you can run Gnome on top of Enlightenment if you wish. Similarly, KDE could be used in conjunction with Sawfish, although KDE has been designed with tighter integration to its own window manager, KWN (you guessed it, the KDE Window Manager).

On an aesthetic level, the easiest way to differentiate between window managers and desktop environments is as follows: the window manager defines how the borders of windows look and operate. The desktop environment dictates how the interior of windows look and function, as the diagram below illustrates.

If you followed our installation guide you should have both Gnome and KDE installed. We suggest you login to both desktop environments and explore a little as you read the following pages. First we'll take a look at Gnome and then KDE.

WINDOW MANAGERS AND DESKTOP ENVIRONMENTS



For those familiar with previous Linux Pocketbooks, Gnome has come quite along way since we last examined it. Now it features an easier to use default configuration, smaller hardware requirements, and the replacement of the default Enlightenment window manager with the smaller and more integrated Sawfish window manager.

Gnome 2.0.1



1 GNOME DESKTOP

Your workspace (also called the Gnome Desktop) incorporates everything you see on screen, including launchers, your panel, and all other items on this list. Double-clicking on an item will open a folder or start an application. Right-clicking enables you to modify the properties for that item.

2 GNOME PANEL

At the bottom of your workspace you'll notice the Gnome Panel, a long bar that looks similar to the Windows taskbar. Rest assured that this is where the similarity ends. The Gnome Panel, as well as providing a quick launching pad for your favourite applications and a list of the programs running, also has many unique features, including nifty panel applets, different panel types, and more. You can explore these options by right-clicking on the panel. Some of the more interesting options include alignment, autohiding, icon sizing, and background images. You can also configure additional panels to sit anywhere on your desktop.

3 MAIN MENU BUTTON

The Gnome footprint is the launching menu for all your applications. It acts much like the Windows Start menu.

4 LOCK BUTTON

If you're leaving your workstation, and don't want anyone using your machine, the Lock Button will come in handy — it immediately activates a password protected screensaver.

5 GNOME CONTROL CENTER

The Gnome Control Center is the central place to configure all aspects of your desktop environment, including window and mouse behaviour, screensavers, and themes for Gnome and Sawfish. You will learn more about customising Gnome in the next section.

6 GNOME TERMINAL

Gnome terminal is a small application which gives you access to the console, the underlying Linux system (you will learn all about the console in Chapter 5).

7 GNOME LAUNCHERS

Gnome launchers are the shortcuts on the Desktop, Panel, Main Menu button, and so on. Gnome supports drag and drop, so it's easy to move launchers between your Main Menu button, Panel, Gnome File Manager and Desktop. Launchers can start applications or open directories in Gnome File Manager.

9 DISK ICONS

Special launchers for quick access to your floppy and CD-ROM drives should appear on your desktop. Right-click on the desktop and select 'rescan devices' if they're not present. These launchers mount (attach) your disk drive to Linux and show its contents in the Gnome File Manager.

10 PANEL APPLETS

Not only does the Gnome Panel provide a handy way of switching between tasks, it can also be extended with applets — small programs that exist entirely on the Panel. These include news tickers, multimedia, system controls, and other small programs with specific functionality you'd like to access on screen. Gnome Pager, Clock, and Tasklist are also applets.

11 GNOME TASKLIST

Tasklist is a Panel applet that displays the applications you currently have open. Try right-clicking on the bar at the end of the Gnome Tasklist and exploring its properties.

12 PAGER

Since many users have a tendency to store their work and particularly important launchers on their desktop, Gnome provides you with four workspaces by default (this number



can be adjusted from within the Gnome Control Center). When your current working environment gets too cluttered, or you'd like to launch a new application off the desktop and would like to minimise all your windows at once, just click on another workspace. The Pager also displays a basic outline of the windows which exist on each workspace, allowing you to see which workspaces contain programs.

13 GNOME FILE MANAGER

The Gnome File Manager, otherwise known as the Gnome Midnight Commander (GMC), is started when you click on a folder on the Desktop or on the Main Menu button under

The Gnome File Manager (also known as Gnome Midnight Commander).

Panel applets, such as those shown here, are small programs that exist entirely on the Panel.





Programs → File Manager. GMC is a Gnome version of Midnight Commander, a very easy-to-use and powerful console (text-based) file manager similar to the old DOS-based Norton Commander. Running `mc` at a console will start Midnight Commander.

GMC is a fairly standard file manager — an expandable tree view on the left, and the contents of the currently selected folder on the right. Gnome File Manager also uses Web-style buttons to navigate and change views on the directory tree.

GMC also has some handy features; among them are the ability to browse FTP sites and strong integration with the Gnome package management software. Take some time to explore Linux with the File Manager, but remember, because Linux doesn't allow regular users do things which could harm your system, you won't have full permission to access many folders outside your home directory.

You can move files using drag and drop, and holding down the Shift key while you drag will copy the file, instead of moving it. Alternatively, you can hold down Alt while you drag the file; a menu will pop up and ask you whether you'd like to copy, move or link to the file when you release the button.

Right-clicking on a file also provides access to a number of options, including the ability to change appropriate properties and manage permissions (access rights).

GMC also has a reasonable level of customisation via its Settings → Preferences menu. Here you can set up a custom view type for displaying the contents of directories, turn on or off the viewing of hidden files, turn warnings for file deletions and other actions on and off, and other handy settings.

KEYBOARD AND MOUSE SHORTCUTS

Gnome's keyboard and mouse shortcuts are easily configurable from within the Control Center, under Sawfish Window Manager → Shortcuts. The table left includes some of Gnome's default keyboard shortcuts.

KEYBOARD SHORTCUTS

SHORTCUT KEY	ACTION
Alt + Tab	Switch between programs
Alt + [F1 through F8]	Switch between workspaces
Shift + Alt + up arrow	Bring a window to the front
Shift + Alt + down arrow	Send a window to the back
Shift + Alt + [F1 through F8]	Send a window to desktop [1 to 8], and switch to that desktop

Customising Gnome 2.0 Gnome and KDE are so configurable it would be a sin not to explore the many aspects of your desktop you can tweak, fiddle and tailor to your heart's content.

Without a doubt the best way to explore how to configure Gnome to your liking is to work your way through the configuration menus trying everything out, and the Gnome Control Center is where all the tweaking takes place. Not only does it let you tailor the look and behaviour of Gnome, but also the Sawfish window manager upon which Gnome depends. The Gnome Panel can also be configured here, but we'll look at that separately after walking you through some of the more interesting features of the Gnome Control Center.

THE GNOME CONTROL CENTER

Start the Control Center by clicking on the toolbox in the Gnome Panel. You'll notice all the facets of Gnome listed in a hierarchically organised tree menu on the left. Let's start from the top and walk through some of the more important menu items.

DESKTOP

The Theme Selector lets you choose from a wide variety of themes you can apply to buttons, menu lists, and so on. There are some fantastic themes to choose from, so explore a little to

see what's available. To set a theme for the window manager, Sawmill, select Appearance under the Sawfish window manager section.

The Window Manager section allows you to switch and configure window managers on the fly. If you're using Sawmill you can configure it directly in the Gnome Control Center, otherwise click 'Run Configuration tool' to configure the selected window manager.

DOCUMENT HANDLERS

Chances are you're not an Emacs person so for now you may want to

The Gnome Control Center.





change it to a WYSIWYG-style editor such as `nedit` or `gedit`.

Mime Types allows you to set file associations for file types. Keep this in mind if you ever want to change the default applications associated with file types.

SAWFISH WINDOW MANAGER

Appearance enables you to choose a Sawfish theme (there are some very nice themes available). Mix and match these with Gnome themes to create your own theme sets.

Focus behaviour is really a matter of preference, but you can configure Sawmill to raise applications when the mouse is rolled over them, only raise them when clicked on or any-thing in between. Shade Hover is a nice feature where, if a window is shaded (that is, 'rolled up' so that just the title bar is visible — which is achieved by double-clicking on the title bar) the window will be auto-matically unshaded when the mouse is over it and then shaded again when the mouse moves on.

Moving and Resizing sets how windows look when they are dragged (opaque, boxed), how they look when resized, the level of window 'snapping' and more.

Shortcuts is useful just to explore the key combinations you can use to perform common tasks. You can also add your own key shortcuts here.

Workspaces sets the number of virtual desktops and how they operate, including edge-flipping options and backgrounds for each desktop.

SESSION

Startup Hint lets you set what type of greeting appears when you login. This defaults to Gnome tips, but for a more interesting greeting click 'Display Fortunes'.

Startup programs is the place to put any programs you want started automatically when you login. This includes, for example, running a screensaver as a background to your desktop (this is very cool). Setting this up is quite simple.

First, click on the Gnome menu and select Run. Type in `xscreensaver-demo`. This will start the `xscreensaver` demo program, a small front end that allows you to browse the screensavers that come with `xscreensaver`. Once you've found one you like take a note of the command that is used to run it. (This will be something like `galaxies -root`, where the 'root' option runs the screensaver in the 'root' window, or in other words your desktop display). From the Gnome Control Center browse to Desktop → Background and turn off 'Use Gnome to set background'. Then scroll down to Startup Programs and click Add. Type in the corresponding command as listed in `xscreensaver` (such as `galaxies -root`) and you're done. Next time you log in, you'll have an ani-mated desktop from which to work. It's important to choose a screensaver that's easy on the eyes and not



too distracting. Remember that the really fancy ones will also chew a lot of your CPU time, so choose carefully. In particular, you might like to check out the coral, bouboule, xflame and xearth screensavers (xearth is actually a stand-alone program, but can be run on the desktop just the same). You can, of course, download more screensavers online and have, for example, The Matrix running on your desktop!

THE GNOME PANEL

The Panel is almost as configurable as the rest of Gnome. The standard settings are self-explanatory but the applets deserve special mention. Gnome applets are small programs that run in the Gnome Panel, and there are plenty to choose from.

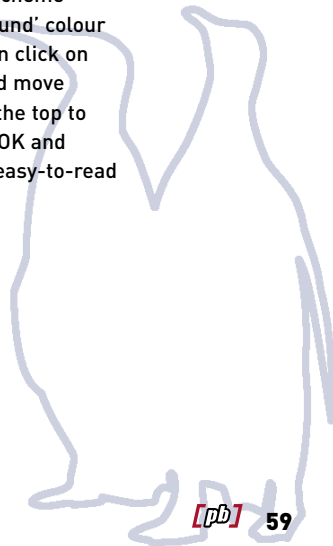
Right-click on the Panel and select Panel → Add to panel → Applet. Some of the more notable ones include the Afterstep Clock, Gnome Weather (which displays the weather in any part of the world while you're connected to the Web), Tick-a-Stat to monitor log files, CPUload, Xeyes, and Webcontrol. The properties for all applets can be changed simply by right-clicking on them. Hold down both mouse buttons over an applet and drag to change its position in the panel.

THE GNOME TERMINAL

Later in this book you'll learn about terminals and how they enable you to issue commands directly to your Linux system. However, we'll show you how to customise the Gnome terminal program in this section as it's a powerful tool that you'll find yourself using a lot in the future.

Start a Gnome terminal by clicking on the icon that resembles a monitor on the Gnome Panel. To make it behave a little more like traditional terminals, click on Settings → Preferences → General and set the following options, according to your taste: Hide menu bar; 'Transparent' and 'Shaded', under the Image tab; and 'Scrollbar position' to 'Hidden' under the Scrolling tab. Lastly, choose 'Custom' for both 'Color Scheme' and 'Foreground/Background' colour under the Colors tab. Then click on the Foreground colour and move the vertical colour bar to the top to achieve pure white. Click OK and you'll have the standard, easy-to-read white-on-black terminal.

The highly configurable Gnome panel.





place to configure all aspects of the K Desktop Environment, including themes, wallpaper, icon sizes, and more.

5 K FILE MANAGER

This link opens your home directory, where you can store all your personal stuff. KDE 1.2's file manager is called KFM while KDE 2.0 and 3.0 has Konqueror.

KFM is a simple graphical file manager, with the added capabilities of displaying basic HTML and thumbnail graphics. KFM behaves in the same way as many popular

file browsers; a directory tree on the left, and the contents of the currently highlighted directory on the right. Files can be dragged and dropped between locations within KFM, as well as on the desktop. Left-clicking on a file launches it, and right-clicking allows you to view and modify properties.

In addition to having all of the features mentioned above, Konqueror is designed to be a full Web browser replacement with support for Netscape Navigator plug-ins, all aspects of modern HTML, Windows and Unix network browsing, FTP, and more. The interface bears resemblance to Internet Explorer (IE) on Windows, with a common set of buttons for navigating Web content and the filesystem. Other IE features, such as the ability to browse Windows shares, or search for words using the address bar, are also implemented.

If you examine the Window menu, you'll find options for hiding the directory tree, a split view (for comparing and manipulating two or more directories on the same screen), a full-screen view (perfect for Web browsing), a terminal window (which echoes your movements in the file manager), and more.

Another noteworthy feature is the Configure Menu (Settings → Configure). Here you can set up file associations (including more than one program per file type), proxies for Web and FTP browsing, search options, and so on. Browse the menus to fully explore Konqueror's capabilities.

6 K MENU

The equivalent of the Windows Start menu, the K menu contains a list of all your applications neatly sorted and accessible, and also allows you to directly browse your filesystem. Unfortunately, KDE 1.2 does not support drag and drop within the K menu, but later versions do.

7 MULTIPLE DESKTOPS

KDE sports a number of desktops for you to use. If you have a selection of different programs open, and you want to start a new application from a desktop icon or get some more screen real estate, just select a free desktop. All your open windows will remain in the taskbar and on the desktop they were originally started, allowing you to select them from any desktop. You can also activate the 'sticky' button (which looks like a drawing pin or circle) to have an application 'stick' across all desktops. Additionally, you can specify different wallpaper for each desktop using the KDE Control Center.

8 DISK ICONS

These icons give you access to your floppy, CD-ROM and other removable drives. When you put in a removable disk, click here to access the contents if a file manager window doesn't automatically open. The icon will display when the disk is mounted. You can right-click to view options, including to unmounting and ejecting the disk.

9 HOME

This links to your home directory, where you can store your personal documents.

10 KONSOLE

Click here to access a plain Linux console. KDE 3.0 users have the option of spicing things up a bit with semi-transparency options and other aesthetic enhancements.

The Konsole and Midnight Commander.



KEYBOARD AND MOUSE SHORTCUTS

SHORTCUT KEY	ACTION
Alt + Tab	Switch between applications
Ctrl + [F1 through F8]	Switch between desktops
Alt + F4	Close the current Window
Tab	Switch object focus
Alt + F2	Bring up a basic 'run command' dialogue
Ctrl + Alt + Esc	Window kill mode — provides a small crosshair you can use to kill any graphical programs that aren't responding.



Customising KDE

The new version of KDE is a powerful and well-integrated desktop environment. In this section we'll take a look at customising KDE 3.0, Mandrake's default desktop environment.

As with Gnome, the best method for exploring how to tailor KDE to your liking is to work your way through the configuration menus. For KDE, the KDE Control Center is the source of all things tweakable. Not only does it let you customise the look and behaviour of the KDE environment, but also the underlying KDE window manager. The KDE Panel can also be configured here, but we'll look at that separately after walking you through some of the more interesting features of the KDE Control Center.

THE KDE CONTROL CENTER

Start the Control Center by clicking on the Control Center icon in the K panel. You'll notice all the facets of KDE listed in a tree menu on the left. Again, we'll start at the top and walk through the menu items, covering some of the more interesting features. At any time you can access helpful information by clicking on the question mark in the window border and then on the feature you'd like to learn more about.



APPLICATIONS

What's Dr Konqi? Much like a Dr. Watson for Windows, it intercepts crashing applications and keeps you informed.

FILE BROWSING

Under File Manager turn off 'underlines' if you want to make Konqueror, the KDE integrated file manager, even remotely usable as a file manager. Even though Konqueror can be used as a fully featured browser, there's no need for underlines when using it in file manager mode.

You'll notice a menu entry for File Associations. Come back here if you should ever find a file you want to associate to a particular application. Take note that, unlike Windows, multiple programs can be associated with a file type. File associations can also be edited by right-clicking on a file in Konqueror and selecting Edit File Type.

HELP

KDE's Help manager is rather nifty. Apart from providing full text searches, which is standard, KDE's Help allows you to browse 'man' pages (you'll learn more about these later) and application help files as well as help on KDE.

INFORMATION

Information doesn't let you tweak much, but does centralise all the standard information about your system. This is a great place to find out, for example, what version of XFree86 you are running, or to make sure sound support is loaded.

LOOK'N'FEEL

This is the place to tweak the look of KDE. This section is rather large, so the highlights are as follows.

Under Actions you can set the focus policy (raising a window by clicking or positioning the mouse over it) and determine how new windows are placed on your desktop.

The Background section lets you do more than just set backgrounds for your desktop. You can configure

a variety of colour gradients instead of bitmaps, set a number of different images to be rotated at a specified interval or even run a program as a background, and you can do any one of these for each of your desktops. If, for example, you're a fan of xearth you can run it as your desktop background by selecting Desktop, setting Mode to 'Background Program', clicking Setup and then choosing XEarth. To ensure it displays properly you then need to click the Wallpaper tab and set Mode to 'No Wallpaper'.

NETWORKING

Windows shares allows you to set up WINS server or browser server addresses to help Konqueror find available Windows shares on a local network using Samba. For the most part you don't need to do anything to start browsing Windows machines on your network, but by providing WINS or browser server addresses you can greatly increase the range of machines Konqueror can find.

PERSONALISATION

This section allows you to set up personal details, such as your email address for use with an email client. There's also a neat feature to remember passwords you use in KDE, with a built-in time limit before it should 'forget' them and prompt again.

SYSTEM

As root you can configure the login manager as well as system variables

such as the timezone. The Menu Editor is a fantastic program that allows you to easily manage the KDE menu system.

WEB BROWSING

This last section configures how Konqueror behaves as a browser. Here you can set how Konqueror will treat cookies, how HTML should be rendered (we recommend you deselect 'Underline links' and enable Java support) as well as manage Netscape plug-ins and proxies. Konqueror is a remarkably good browser, and more stable than Netscape Navigator, despite still being under development.

THE K PANEL

There's plenty to play with here so we'll just show you some of the highlights to get you started.

Right-click on the panel and select Configure → Settings. From here you can set the panel size, hiding options, the location of the panel and other useful features. One way to brighten the panel is to click on the Buttons tab and select 'Enable background tiles' for all button types. If you don't like the default tile set, you can browse from the list available.

As with Gnome you can also add applets to the K panel by right-clicking on the panel and selecting Add → Applet.

THE KONSOLE

To tailor the KDE terminal program, start a Konsole by clicking on the icon that resembles a monitor on the K Panel. To make it behave a little more like traditional terminals, right-click anywhere in the Konsole and select Schema → White on Black. Play around with the other schemes, but chances are you'll settle on white on black as it's the easiest to read.

If you'd like more space in the window (and trust us, you will want more space later on) you can remove the scroll bar by choosing Scrollbar → Hide. You can return the scrollbar at any time if you ever need to scroll back by reversing this option. To free up a little more space you can also deselect 'Show Menubar', as you've already discovered all the options in the menu bar are available by just right-clicking in the Konsole. Lastly, set the font to be a size which leaves plenty of room for information onscreen without being too hard to read.

The toolbar at the bottom is a fantastic idea and allows you to start up multiple terminals, and even the Midnight Commander file manager, in the one window. Simply right-click on it to set options such as position and size. When you're happy with the settings, select Save Options.



Using Linux

CONNECTING TO THE NET 68

INSTALLED APPLICATIONS 71

CONFIGURING LINUX 75

UPDATING LINUX 84



Connecting to the Net

connecting to the Internet with Linux is far easier than it has been in previous distributions. Automated tools now allow you to set up an Internet connection with a minimum of fuss.

Linux was built around, and using, the Internet. As a result it supports a wide variety of Internet connection methods, including dialup modems, ISDN, ADSL, cable, and more. Linux also supports advanced networking features such as firewalling and masquerading — the ability to allow others to access the Internet through your machine.

CONNECTION DETAILS

You'll only need to know a few details about your connection to get it working under Linux. If you are dual booting with Windows, you can see what you've set up for your Windows connections in Dial Up Networking. Specifically, you need to know:

- ☞ Your ISP's phone number
- ☞ Your username
- ☞ Your password

This should be all the information you need. In the event that things don't work the first time (which is very rare), knowing some of the extra details on your system can help. Specifically:

- ☞ Which serial port your modem is attached to on your computer. Windows terminology uses the terms 'COM1' through to 'COM4'. In Linux, these same ports are referenced as /dev/ttyS0 through to /dev/ttyS3.
- ☞ Your ISP's primary DNS server (and secondary DNS server, if one exists)

USING A MODEM WITH RED HAT 8.0

The Red Hat PPP tool (otherwise known as `rp3`) is one of the easiest dialup Net connection tools for any OS. As with Windows, the Red Hat PPP tool will autodetect your modem and automate some of the connection details. But unlike Windows you won't need to install any extra drivers or networking capabilities, and you won't have to reboot.

To start, simply click the Dialup Configuration icon on your desktop and enter the root password when prompted. If a modem wasn't set up during installation, you'll

be told that you need to configure one. Click Next and let the program search for a modem (make sure your modem is switched on). All the details about the modem will be automatically detected.

If you're happy with the results, click Next, otherwise you can manually add a modem before continuing.

You'll be asked to provide the phone number and a name for the connection. Type in the name of your ISP and its phone number, then click Next. On the following screen fill out your account details (name and password) — and remember these are case sensitive.

The final screen will ask you if you need any special options. Click Next and you'll be presented with a summary of the options you've chosen. Click Next again to finish the Wizard, and a new entry will be made in the Internet Connections dialer.

If you have problems connecting later, turn on the Debug option to view a connection window that will display the progress of connecting to the ISP. Using the Advanced tab on the Internet Connections window to provide extra detail might also help.

Now that you've set up your connection, you can launch it (as a regular user) by clicking Programs → Internet → RH PPP Dialer in your Gnome menu.

USING A MODEM WITH MANDRAKE LINUX 9.0

Mandrake Linux uses the KDE PPP connection program, KPPP, to auto-



mate connection to the Internet. Simply click on the 'Internet' icon on your desktop to launch KPPP. Once started, click on Setup to configure a new connection. Click New and select Dialog setup. Type in a connection name and then click Add to add a number for the connection. Explore the other tabbed areas if you like, but everything should be configured for you.

After creating the entry, it should appear on the main KPPP screen. At the top of this screen there are a number of tabbed areas, the first of which is Device. Click here and then select which serial port your modem is connected to. As mentioned earlier, this will be one of ttyS0 through to ttyS3. You will see an entry for /dev/modem, but this won't be configured unless you used DrakNet to set up your modem during installation. Lastly, click on the Misc tab and select 'Dock into panel' to have KPPP minimize to the panel after connecting.

The Red Hat PPP connection program.



With your connection highlighted, type in your username and password then click 'Connect' to do so.

USING A CABLE MODEM WITH RED HAT AND MANDRAKE

Telstra and Optus cable modem services work exceptionally well with Linux. Both use DHCP to automatically assign IP addresses to clients and attach to the network. However, both providers require additional steps to authenticate with their servers.

KPPP, the KDE PPP connection program.

TELSTRA BIGPOND ADVANCE

Telstra users must first install the `bpalgin` program from www.trontech.net/bpalgin. Simply grab the package named `bpalgin02.0-1.i386.rpm` and install it with: `rpm -ivh bpalgin02.0-1.i386.rpm`.

The next step is to configure DHCP if you don't already have it running. The easiest way to do this is to use `linuxconf`. `su as root`, start `linuxconf` and click through to `Networking` → `Client tasks` → `Hostname and IP network devices`. Then click on the adaptor you are using and select DHCP. Lastly edit the `/etc/bpalgin.conf` file and modify it to include your account details. To start the service, simply run:

```
service bpalgin start
```

There's no need to run this when you reboot, as `bpalgin` will start automatically.

OPTUS@HOME

Optus@Home users do not need any additional software, but must send their hostname (on the inside cover of your Optus@Home CD) as part of the DHCP negotiations. You can do this with the DHCP Client Daemon (`dhcpcd`), using the following command when logged in as root:

```
/sbin/dhcpcd -d -h [hostname] eth0
```

If the above works, you can automate the connection process by adding this to the end of `/etc/rc.d/init.d` to load as part of your startup sequence.

For more information about cable Internet services under Linux visit Cable Users Australia at www.cableusers.com.au.

Installed applications

You've installed a new Linux system, customised it to your liking, connected to the Web, and the world is waiting at your feet. Now what?

It's time to explore the many programs that come preinstalled with Linux! One of the reasons Linux is distributed on two or more CDs these days is that just about any program you'll ever need to use is provided for you in the box as it were. Everything from word processors and server programs through to development tools and multimedia applications.

In this section we'll highlight just a few of the preinstalled programs supplied with Red Hat 8.0 and Mandrake 9.0. Some of these you'll find buried in the menu system under Gnome and KDE, while others can simply be run from the command line (program name provided in brackets below). Unless otherwise stated these programs can be found under both distributions. If the program isn't already installed, you can install it yourself by browsing the distribution CDs using a package manager (see page 113).

GRAPHICS

GQVIEW (`gqview`)

GQView is an image viewing application similar to ACDSee or Infranview on Windows.

THE GIMP (`gimp`)

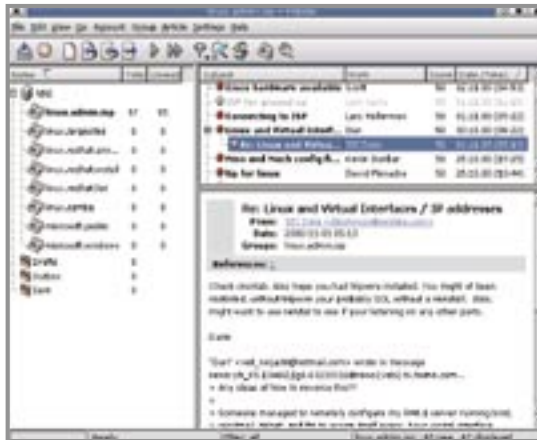
The GNU Image Manipulation Program (hence The GIMP) is one of the oldest and most established Gnome-based Linux applications. It's a fully featured photo manipulation suite with layers, filter effects, a powerful scripting language, vector-based text, and much more.

ELECTRIC EYES (`ee`)

From the author of the Enlightenment window manager comes a small, simple program for viewing images and lists of images.

The GNU Image Manipulation Program.





FILE TRANSFER

GFTP (gftp)

A simple drag-and-drop FTP program that has a stack of features, including bookmarks, session logging and troubleshooting tools.

NcFTP (ncftp)

NcFTP is a traditional command line FTP client, but don't let that put you off. NcFTP offers bookmarking, resumed downloads, autologin, command completion, and much more. Easily one of the best FTP clients for any platform.

Knode is an elegant KDE 2.0-based newsreader.

MIDNIGHT COMMANDER (mc)

Midnight Commander is a clone of the popular Norton Commander for DOS. Certainly one of the most intuitive, quick and feature-rich file managers for Linux. Managing your filesystem — whether you're in Gnome or KDE, or at the console — is a simple affair with Midnight Commander.

ARK (ark)

KDE's archiving tool, Ark, provides a simple way of creating, modifying and extracting a wide variety of archives, including tar, GNU zip, BZip2, Windows Zip, Zoo and RAR.

COMMUNICATION

LICQ (licq)

LICQ is an ICQ client for Linux. It's fully compatible with the Mirabilis ICQ network and with all the features of Windows ICQ (including the ability to register new ICQ users). LICQ also supports a variety of plug-in modules and skins.

XCHAT (xchat)

An excellent IRC client for Linux, XCHAT provides all the basic chat features but includes a wide variety of less common functions such as scripting, user lookup, broadcasting and some helpful security tools.

EVERYBUDDY (everybuddy) [MANDRAKE 9.0]

Most instant messaging applications only allow you to use one of the four major competing instant messaging networks: Mirabilis ICQ, AOL Instant Messenger,

Microsoft MSN or Yahoo's IM Network. Everybuddy is designed to allow open communication between all these networks, allowing you to log on to all four simultaneously and interact with users in a seamless manner.

KMAIL (kmail)

Kmail is, as you'd expect, a mail client for KDE. While KDE 1.2 provides simple POP3 email capabilities, KDE 2.0 and later provides much more functionality, including IMAP, PGP-based encryption, the use of external editors, a shared KDE-wide address book and spam filtering.

KNODE (knode) [MANDRAKE 7.2]

A KDE 2.0-based news client which is good looking, easy to use and comes with a number of excellent features, including offline news browsing, kmail integration, searching, inline attachment display and content filtering.

MULTIMEDIA

X-CD-ROAST (xcdroast)

If you have a CD burner, X-CD-Roast is your one-stop CD-burning shop. It provides automated setup of your burner, on the fly or image-based copying, audio CDs with track names grabbed from the Internet, excellent hardware support, and more.

GRIP (grip)

A Gnome-based all-in-one CD to MP3 solution. Grip is a simple yet powerful application to build your MP3 collection with a minimum of fuss. Behind the scenes, Grip can use



Grip is an all-in-one solution for MP3 creation.

a number of different command-line MP3 (and other format) encoders and ripping programs, allowing audio-philosophers to pick their favourite open or closed source tools for the job. Grip supports grabbing track names from the Web, multiprocessor encoding and SQL databases, and can also act as a regular Internet-aware CD player.

X MULTIMEDIA SYSTEM (xmms)

XMMS is the Linux version of Winamp, supporting the same range of features, skins and playback quality of its Windows counterpart. This is all you need for high-quality MP3 playback.

OFFICE

ABIWORD (abiword)

A basic word processor application that is fast and elegant to use. AbiWord supports Microsoft Word input and output as well as a variety of other features.

**GNUMERIC** (gnumeric)

A Gnome-based spreadsheet application that originally started out as a clone of Microsoft Excel and has since become the premiere Linux spreadsheet program.

KOFFICE [MANDRAKE 9.0]

Part of standard KDE 3.0 installations, Koffice is a fully functional office suite with a word processor and desktop publishing program (kword), spreadsheet (kspread), presentation application (kpresenter) and vector drawing tool (killustrator).

Red Hat's Gnome-lokkit utility helps you set up a basic personal firewall.

MANAGEMENT**USERDRAKE** (userdrake) **[MANDRAKE 9.0]**

An extremely simple method for adding, editing and deleting user and group accounts on your system.

KDE SYSTEM GUARD (ksysguard) **[MANDRAKE 9.0]**

KDE 3.0 comes with a flexible tool to measure and monitor system performance. The simple interface allows you to graph performance data for CPU, memory, swap, network, disk activity and a whole lot more. System Guard can also be used to monitor multiple machines over a network.

GNOME-LOKKIT (gnome-lokkit) **[RED HAT 8.0]**

A Wizard-based firewalling program to help prevent potential crackers from breaking in to your system. Gnome-lokkit uses the firewalling capabilities of IPchains built into the Linux kernel, but adds a simple question-and-answer-style interface to determine exactly the type of incoming connections you like to receive.

GKRELLM (gkrellm)

This nifty program will display and update just about any aspect of your system you care to monitor — from CPU, memory and swap utilisation through to network throughput and even if you have new mail. A variety of skins can be used to help Gkrellm blend into the theme of your desktop.

Be sure to check out the Pocketbook cover CDs which highlight some of these programs in addition to a selection of popular programs not available in either distribution.

Configuring Linux

Linux is more than the beautiful interface of Gnome or KDE. Its flexibility allows you to tailor Linux to almost any task, and flexibility is achieved through configuration.

Although every single aspect of Linux can be configured through the command line, much of it can also be configured through GUI tools that help automate the task of configuring Linux. The following guide will walk you through the various tools you can use and the settings they configure. If you're planning to use Linux as a desktop OS, you won't need to change much — if anything — but it helps to know the tools that are available and when you might need to use them.

THE MAGIC OF SU

When it comes to system-wide settings, only the root account can make changes. Many of the configuration tools listed here will prompt you for the root password before allowing you to launch them. There may be some tools which won't prompt for the root password and will, instead, politely inform you that you don't have access to run that program. While you could log out and then login again as root, a much easier way is to login as root within a terminal — the Linux equivalent of the command line. This allows you to run root-privileged programs from your user account.

Terminals are covered in-depth in the next chapter, but for now simply click on the terminal icon (on the panel) to launch a terminal and then login as the root user by typing:

```
su -
```

You will be prompted for the root password. The `su -` command actually allows you to switch to any specified user, but defaults to the root account. From here you can issue commands just as you would if you were using the terminal in your user account, except that while logged in as root you will have full access to the system.

You can exit out of the root account at any time by typing `exit`. This returns you to the terminal where you are, by default, logged in as your user account. Typing `exit` again will log you out of the terminal, and thereby close it.

CONFIGURING RED HAT 7.0

Red Hat's configuration tools have changed significantly in recent releases. The once trusty (and still present) control-panel tool we examined in the previous Linux



Pocketbook has been superseded by the all-purpose configuration masterpiece that is linuxconf (covered later). Linuxconf didn't last long, though, and was removed with the advent of Bluecurve. Setup is still the favourite for Red Hat.

SETUP (setup)

Each menu item within setup is actually a separate program which can be launched independent of setup. We'll highlight just the important tools, in particular those whose functions might not be immediately obvious.

Authentication Configuration (authconfig)

This tool can be used to choose different methods of storing user account information, and checking the credentials of those attempting to log on to your system. Chances are you won't need to know about authconfig, since all the common security precautions are activated, but it's helpful to explore in order to gain an insight into Linux authentication. More on the authentication services can be found in *The Advanced Linux Pocketbook*.

System Services (ntsysv)

A simple menu to configure the services that run automatically at boot time. Pressing F1 will give a description of each service. Starting services is covered later in this section.

Sound Card Configuration (sndconfig)

Unlike Mandrake, Red Hat doesn't configure your soundcard during installation, so you may want to set this up now. Configuring sound in Red Hat is easy — if you have a supported soundcard. Popular brands include the Sound Blaster Live series, ESS and most SoundBlaster-compatible cards.

The sound utility will attempt to detect your soundcard even if it isn't directly supported. If you're lucky you should hear Linus Torvalds' voice telling you the correct way of pronouncing Linux (it's 'Lyn-ooks', by the way). If not, then there's still a good chance you can get sound working (see page 83).

X Configuration (Xconfigurator)

Red Hat has made substantial changes to this tool since the previous Pocketbook. Xconfigurator will now autodetect a far wider range of graphics cards and monitors, to make sure you not only get full support for your graphics card, but also the maximum refresh rate from your monitor.

In the rare event that Xconfigurator has trouble setting up your video card, you may want to try XF86Setup, a similar utility produced by the XFree86 team. Take note that these tools are designed for XFree86 3.x and may not work with XFree86 4.2.

CONFIGURING MANDRAKE 9.0

Since Mandrake is built on Red Hat, it features many of the same configuration programs (such as linuxconf) to configure a Mandrake system.

Mandrake has also written its own extensive and user-friendly configuration tools, all centralised through the control panel-style DrakConf.

DRAKCONF (DrakConf)

DrakConf acts as a simple interface to a number of system administration utilities, all of which require root access to run. Simply click the icon on your desktop, enter the root password when prompted, and you'll be able to access all the utilities as root.

Menu Editor (menudrake)

A menu editor that allows you to simultaneously modify the menus for KDE, Gnome and other window managers on your system. This ensures that applications appear consistently in the same menu categories no matter which window manager or environment is used.

Package Manager (rpmrake)

This nifty program is a powerful package manager (for more on packages see page 113). RPMDrake lists the available uninstalled packages on both Mandrake CD-ROMs by default, but you can point it to any directory (such as a download directory) for new programs.

Startup Services (drakxservices)

Configures programs and services

to run at bootup. This program performs the same task as ntsysv and both are interchangeable. Services are covered later in this section.

Font Manager (drakfont)

An excellent utility for configuring and installing TrueType fonts, including the ability to locate and install Windows fonts found in any Windows installations on your system (if you have a dual-boot system with the Windows partition mounted). We highly recommend doing this; the difference Windows TrueType fonts can make in your Linux applications — including Netscape — has to be seen to be appreciated.

Network and Internet Configuration (draknet)

A utility for connecting to the net via modem, ADSL, cable, and more.

DrakConf provides a centralised interface for all your system configuration tools.



**Security Level (draksec)**

Configures Mandrake's predetermined security levels. Medium-level security should suffice for most desktop machines, and high should be adequate for servers, though this is certainly no substitute for a professional security audit.

Mandrake Update (MandrakeUpdate)

A utility for downloading bug fixes, program updates and security patches from MandrakeSoft or a local mirror. For Australia, this will likely mean the AARNET file servers in Brisbane. MandrakeUpdate is covered at the end of this chapter.

Live Update (livedrake)

This is easy to confuse with MandrakeUpdate as a means to update packages, but livedrake actually lets you install packages from Mandrake installation CDs, thus providing a method to upgrade to newer versions without rebooting.

Change Screen Resolution (XFdrake)

Provides a quick menu to configure your X resolution and colour depth settings.

Hardware Configuration (harddrake)

An excellent diagnostic tool for Linux which also provides some configuration options. Although harddrake is only in beta stage, it is quite stable. The application is styled after the Device Manager from Windows 9.x and 2000, and aims to provide a similar interface for configuring hardware on Linux.

Internet Connection Sharing (draknet)

A simple yet excellent tool that allows you to set up Internet connection sharing with the click of a button. This program simply takes advantage of IPChains rules for masquerading, but it's a welcome addition by Mandrake to make sharing Internet connections so easy.

Tools for Booting (drakboot)

Drakboot is a handy interface for creating boot disks and formatting floppy disks. It also offers basic configuration for your boot loader (LILO or GRUB).

LINUXCONF

Linuxconf is both Red Hat and Mandrake's primary configuration tool for most administrative tasks. It is a plug-in based expandable system, and easily one of the most comprehensive all-in-one configuration tools for Linux.

Red Hat users can start Linuxconf by suing as root and then running linuxconf; Mandrake users can start Linuxconf from within DrakConf (click on System Configuration).

Because Linuxconf aims to be a central repository to configure all aspects of a Linux system, it's far too large to explore in this book, but we do recommend you explore it for yourself to discover the many facets of Linux you can configure. Many of the options you will find form a sizeable chunk of the type of services and settings that will be covered in *The Advanced Linux Pocketbook*. Some areas you may want to explore include:

Networking → Client Tasks → Basic Host Information

Configure your machine's hostname, domain name, IP address and multiple network adaptors.

Networking → Server Tasks

This subsection enables you to configure the Apache Web server, anonymous FTP and even Samba for Windows networking.

Miscellaneous → Modem

Set up a /dev/modem link.

Boot configuration → LILO

Helps you manage LILO entries.

Control → Date and Time

Sets the system's date and time.

Status → Logs

Lets you view kernel, boot and system messages.

An alternative to Linuxconf for Mandrake 7.2 is the rather impres-

sive Webmin tool. Webmin allows you to manage a machine entirely through a browser, and thus over a network or the Internet. Webmin isn't as comprehensive as Linuxconf but does provide a configurable front end to just about any service or feature you will ever need — you can even partition drives, setup RAID arrays, manage servers, upgrade software (and burn CDs with the latest version), all through the convenience of a browser.

To check out Webmin, simply make sure the 'webmin' service is started (see the section on 'services', coming up), open your browser, and connect to your machine (the local-host) with: <https://127.0.0.1:10000>

CONFIGURING SERVICES

Services are programs that run in the background to support various functions such as Web serving, error logging, print spooling and so on. By default your Linux installation will be running a variety of services, many of which are required for proper system operation. However you can improve boot time and save system resources by disabling services you don't need or intend to use.

Services (also known as daemons) are covered from a security perspective in the security section on page 145 and the use of services themselves will be covered in *The Advanced Linux Pocketbook*. For now we'll just explain how to disable unnecessary services to free up system resources.



SELECTING STARTUP SERVICES

Both Red Hat and Mandrake provide a simple console based service selection tool called ntsysv. For this exercise start a terminal (click on the terminal icon from the panel), use the su command to login as root, and then run ntsysv.

You should be presented with a simple interface that lists all the services for your system with an '*' next to those that are scheduled to start at boot time. The number of services listed depends on your system configuration and what you've have installed. Use Tab and the cursor keys to move around, and the spacebar to select/deselect. Pressing F1 will give a brief description of the service.

The table opposite lists the common services (this is by no means all the services) and what they do.

If you simply plan to use your machine as a desktop system connected to the Internet, and don't intend to run any servers, you can safely disable the following services:

arpwatch	rawdevices	snmpd
netfs	innd	squid
identd	ipop3	sshd
nfs	named	ybind
nfslock	postgresql	postfix
portmap	routed	

Depending on your system's configuration and what you plan to do with your machine, you can also disable the following:

pcmcia	lpd	proftpd
usb	telnet	httpd
linuxconf-web	wu-ftpd	webmin

If you're not sure whether to disable a service or not, leave it enabled. Some services are essential to your system's well being, and disabling them could effectively disable your system.

When you've finished running through the list, hit 'OK' and your changes will be saved. Next time you reboot, only those services you've selected will be started.

PRINTING

Linux can connect to a wide range of locally attached printers as well as remote printers shared from machines running Windows, Netware and Unix-like systems. Setting up printers under Linux is easy, and Linux will even detect where your printer is connected.

COMMON LINUX SERVICES

SERVICE	TASK	SERVICE	TASK
atd	Provides a way of scheduling tasks to be performed in the future.	nfslock	File locking for NFS
amd	Handles the automatic mounting of filesystems.	pcmcia	Services for PC Card style expansion cards.
anacron	Runs scheduled tasks [cron jobs] that were missed previously due to system downtime.	portmap	Port mapping for NFS and NIS.
apmd	Power management server	postfix	A mail server.
arpwatch	Resolves network card MAC addresses to IP addresses.	postgresql	A SQL database server.
autofs	Mounts filesystems on demand.	pppoe	Starts any PPP Over Ethernet connections, used to connect to certain ADSL providers.
crond	Runs maintenance tasks scheduled to be run periodically.	proftpd	An FTP server.
cups	The Common Unix Printing System	random	Provides quality random number generation.
gpm	Gives mouse access in text mode Linux programs.	rawdevices	Allows applications, such as databases, to use raw devices.
httpd	The Apache Web Server.	routed	Provides routing services.
identd	A server to allow another machine to tell which user is running a particular network connection.	sendmail	The sendmail email server.
inetd	The Internet super server. Manages Internet related services.	sound	Starts sound support.
innd	The news server.	smb	Starts Samba server support.
ipop3	Mail delivery server for POP3.	snmpd	A Simple Network Management Protocol service.
ipchains	Firewalling and Proxy Serving.	squid	A Web cache and proxy server, which stores frequently used information from Web sites.
keyrotate	Sets up the keyboard type.	sshd	Encrypted remote access to command line and GUI programs.
linuxconf	Performs Linuxconf tasks.	swat	The Web-based administration tool for Samba.
lpd	Line Printer Daemon.	syslog	Records important events in the history of your system.
named	DNS server.	telnet	Provides remote telnet access.
netfs	Mounts any Unix, Windows, and Netware file shares that you have configured.	usb	Starts USB support.
network	Starts your network interfaces. This must be turned on for day-to-day use, even if you don't have a network card.	ybind	Prerequisite for NIS.
nfs	The Network File System. Allows file sharing with other Linux machines on your local network.	webmin	Web-based system configuration tool for remote administration.
		wu-ftpd	The popular FTP server.
		xf86	The X Font Server. Provides TrueType, type 1, and other fonts to your system.
		xinetd	The new version of inetd.



PRINTING IN RED HAT LINUX 8.0

Red Hat uses the traditional Line Printer Daemon (LPD) to manage print services. Configuring Red Hat for use with a local or remote printer is achieved using the Red Hat Printer Tool. To run it, first su as root and then run `printtool` to get started. The program is self-explanatory, so the only point to note is that if you can't find your printer model from the large list of supported printers, try selecting an older version of your model or the generic Postscript driver — this should work fine for all printers.

PRINTING IN MANDRAKE 9.0

Mandrake makes use of the all new Common Unix Printing System (CUPS). The new system is an easier and more flexible printing system than LPD. It's even possible to configure and manage printers through a browser, much like using Webmin (mentioned earlier) to manage a system.

To start CUPS simply run `DrakConf` and then click on the Printer Configuration Tool. As with Red Hat, the program is self-explanatory and the only point to note is that if you can't find a driver for your printer simply select an older version of your printer model or choose the generic Postscript driver — which should work fine for all printers. You may also want to give the printer a more informative name than 'lp0'.

To make use of the Web-based interface, make sure the 'cups' service is selected to start at bootup and then run the CUPS Web Administration tool from the Configuration → Printers menu.

The CUPS
Web



SOUND ON LINUX

The Linux sound architecture is currently undergoing a revamp. If your soundcard wasn't detected and configured during installation, all is not lost.

Linux previously relied upon a system known as the Open Sound System (OSS) to support soundcards. The company who designed the system, 4Front Technologies, open sourced their sound infrastructure and a small set of drivers, but required purchase of the fully featured drivers to support the full range of soundcards available. As a result, an open source system called the Advanced Linux Sound

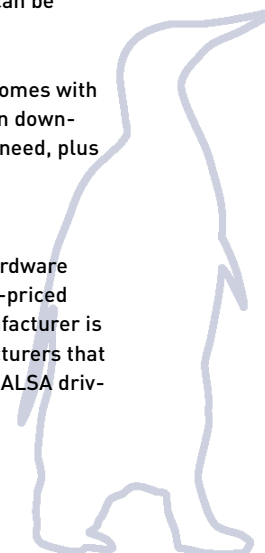


With sound support installed you'll be grooving away to your MP3 collection in no time.

Architecture (ALSA) has been under development to provide the same services as well as being backward compatible with OSS. Although the open source portions of OSS is the more popular system (Red Hat uses OSS), ALSA is slowly becoming more widely accepted (Mandrake uses ALSA).

If your soundcard isn't currently supported by either distribution, there are four options available to you:

- ❏ Check for native Linux soundcard drivers online at the manufacturer's Web site. The Aureal Vortex-based range of cards is one example of a soundcard chipset not supported out of the box but where native drivers can be downloaded. Aureal has since gone under, but Linux programmers have continued developing the drivers Aureal started and the results can be found at sourceforge.net/projects/aureal.
- ❏ Investigate an ALSA driver if you're using Red Hat — ALSA now comes with support for a much wider variety of soundcards than OSS. You can download packages with the libraries, drivers and setup software you need, plus instructions for setting up ALSA on your system. Visit www.alsa-project.org.
- ❏ Another option is to buy a new soundcard that features on the Hardware Compatibility List for your distribution. A reasonable quality, low-priced soundcard which comes with a set of drivers written by the manufacturer is a good idea, and it's also a way to show your support for manufacturers that take the time to produce Linux drivers. A Creative Vibra 128 with ALSA drivers is a good choice.
- ❏ Purchase a commercial OSS driver from 4Front Technologies at www.opensound.com for a small fee.



Updating Linux

Keeping your system up-to-date not only increases its functionality as newer and better versions of programs are released, but it's also important for maintaining a stable and secure system.

Software isn't made perfect. Developers continually improve upon their programs and previously undiscovered bugs are found and fixed. New software regularly appears for Linux, not only because much of the operating system is always being enhanced, but also because the open source community is extremely active at responding to issues such as security. Checking for newer software is an important part of running any operating system, and Linux's smart packaging system and updating utilities make it very easy to do so.

And besides, the performance freak in all of us just loves getting hold of the latest version of our programs, right?

UPDATING RED HAT

The Red Hat Update Agent (known on the command line as `up2date`) will fetch official update packages directly from Red Hat. These packages include enhancements, bug fixes and security updates. Once you've registered with the Red Hat Network (www.redhat.com/network), Update Agent connects to Red Hat and creates a list of available upgrades for your system. Browse through the list, read the descriptions and place a box next to the software you'd like to download. When you click Next, Red Hat will work out any necessary dependencies as well as download and install the packages for you.

UPDATING MANDRAKE

The Mandrake Update (`MandrakeUpdate`) utility, accessible through `DrakConf`, works in much the same way as the Red Hat Update Agent. Mandrake Update provides a simple way to download and install official Mandrake Linux updates from various mirror sites around the globe.

When you first run Mandrake Update you should explore the preferences and, while connected to the Web, update the list of mirror sites. When you want to retrieve a list of updated packages simply click on the big 'update' button and then, after the program connects and downloads a list of packages that apply to your system, select the packages you'd like to install and click 'Do updates' to download and install them.



Learning Linux

LEARNING LINUX 86

TOOLS OF THE TRADE 95

PERMISSIONS 102

SHELL OPERATION 106

MOUNTING FILESYSTEMS 110

INSTALLING NEW APPLICATIONS 113

COMPILING THE KERNEL 117

SURVIVING IN LINUX 123

CONFIGURATION FILES 129



Learning Linux

Linux hasn't always had a powerful graphical environment through which users can interact. In fact, most server systems don't run a GUI at all.


Linux is flexible enough to let you decide whether or not you need to run a resource-heavy graphical environment. Once upon a time Linux would, by default, boot into what's known as the Linux console.

TERMINALS

The console, also called a terminal, is not dissimilar to the humble DOS prompt. If you're an experienced user, you may remember that DOS uses a command shell called COMMAND.COM to process commands, and Linux is no different. The default shell for both Red Hat and Mandrake is bash (the Bourne Again Shell). There are many different shells you can use, each with different features, but they all perform essentially the same function — allowing you to control Linux through the use of simple commands.

Terminals can also be used under X and are known generically as xterms. They are, essentially, a window to the underlying Linux system. Despite the proliferation of graphical tools in Linux, it's an operating system best controlled, and more easily controlled, through a terminal. As such, we'll introduce you to many components of your new Linux system using terminals and the humble command line.

First, start by opening a terminal on your desktop. You will find that you have a choice in the type of terminal you want to use — a simple xterm, the more feature-rich Gnome Term for Gnome or the Konsole for KDE are just a few. They perform the same task, but each comes with different features to make using a terminal easier and more efficient (or, in the case of Eterm for Enlightenment, to make using a terminal a whole lot sexier).



The quintessential terminal

You will find a list of terminals under the 'Terminal' menu from the Gnome and KDE menus, or you can click on the button that looks like a small screen from the panel.

Once started you should be presented with a terminal window that displays your login, the machine name and your current directory. Now we'll show you some commands you can use to get around Linux.

BASIC COMMANDS

Like any operating system, there are various commands and procedures you can use to navigate the filesystem. You will need to know at least a few basic Linux commands to help you get around:

Command	Purpose
cd	Change directory
ls	List a directory
cp	Copy a file
rm	Remove (delete)
cat	View files
exit	Exit the terminal
Tab key	Command completion

Not that hard really, is it? Try `cd /` to change to the root directory and then type `ls` to see what's in it. Note how the output of the command is very basic. The DOS command `dir` also works instead of `ls`. To get a directory listing that shows more information, try `ls -la`. As you might guess, there are plenty of optional switches for

commands like `ls`, and later on we'll show you how to better tailor this and many other commands.

Now would be a good time to point out two distinctions that separate Linux from other PC-based operating systems.

DIRECTORIES

In Linux, directories use the forward slash `/` not the backslash `\` that you may find in other operating systems. It may look odd at first, but you will get the hang of it quickly. The very top level of your filesystem is the root, represented by `/`. All directories branch off from the root.

CASE SENSITIVITY

Linux is a case-sensitive operating system, which gives vast scope for filenames and directories. However, it can sometimes lead to confusion when you change directories or run a program. If you try to do either and it doesn't work, and you're certain spelling is correct, check to make sure there isn't a capitalised letter in there. If there is, you'll need to spell it out exactly the same way.

MAKING NAVIGATION EASY

As you might imagine, trying to navigate long, case-sensitive filenames could be tiring. Enter the magic of the Tab key. No matter where you are in your filesystem, the Tab key will do its best to finish any command you type by matching the contents of the current directory with the letters you type at the command prompt.



Everything in Linux exists on the file-system tree. Although this might sound odd, it's a very powerful model.

For example, if you want to change to a directory called 'Documents', all you need to type is `cd D` and then press Tab. The directory name will be completed for you, then all you need to do is press Enter. If more than one

directory or file matches what you type in, you will need to supply a few more letters to help make the match. If you're not sure of the exact letters needed to complete a match, just press Tab twice and you will be presented with a list of files or directories that match the letters typed in so far.

Using Tab should become second nature to you; after a while you'll find you will be able to rapidly navigate your filesystem regardless of how many levels down a directory is or the length and case of its name. With Tab you don't have to worry about case sensitivity unless it's the first letter. Once you've practised using it you'll wonder how you ever managed to live without it.

THE STRUCTURE OF LINUX

Everything in Linux exists on the filesystem tree. Although this might sound odd at first, it's actually a very powerful model. Apart from the files and directories on your system, you will also find the following aspects of your machine in the filesystem hierarchy:

Devices — /dev

Devices such as hard drives, mice, modems and printers are all represented by a file on your system. Remember the discussion about how Linux names partitions? (See page 36.) Those partition names can be seen in the /dev directory. Operations that you can perform on a file will generally work with devices. For example, you can play a sound file by copying it to the /dev/audio device, or print a text file by copying it to the /dev/prn printer device.

Mounts — /mnt

Mounting is the process by which you attach external filesystems to your directory tree. For instance, you might want to access your Windows partition so you can copy files between your Windows partition and your Linux partition. You can mount the Windows partition onto your filesystem so that the partition becomes just another directory. If you mounted your Windows C: drive as /mnt/cdrive, then typing `cd /mnt/cdrive` will change you over to the root directory of your Windows partition. You could change directly to the Windows directory itself by typing `cd /mnt/cdrive/windows`. You can learn more about mounting on page 110.

System — /proc

The /proc directory stores real time information about your PC, such as memory resources, CPU information, network traffic, and so on. If you change to the /proc directory and type:

```
cat meminfo
```

you will see a summary of your memory usage at the time you typed the command. Similarly, `cat cpuinfo` will show you information about your CPU.

Although you can change to the /proc directory and view the contents of files found there, the files themselves don't actually exist. The files represent various aspects of your system, defined and stored in a logical way in the filesystem hierarchy. You may notice, for example, a rather large file called `kcore`; this is the entire contents of your system's memory.

Since it exists as a file, you could (if you were insane enough) copy the contents to your printer device to print it, to the sound device if you have a liking for binary noise, or to your console if you wanted to view it. This sort of flexibility is typical of the design philosophy behind Linux.

MANAGING PROCESSES

In a terminal, type:

```
ps
```

You should see a list scroll by. Use your mouse to resize the terminal if it's too small so you can see

more information. The `ps` (Process Status) command shows you the processes (or programs) running on your system. In its default state, `ps` will only show processes that you, as a user, have started. However, Linux is always running a variety of background tasks. Type the following:

```
ps aux
```

The letter 'a' represents 'all users', the letter 'u' stands for 'user format', and 'x' tells `ps` to display all processes running on the system. By using these switches we can display more information about each process. The list shows how much memory and what percentage of CPU time each process is using. You can use the information provided in this list to manage your system. For instance, although you should rarely find a need, you may want to kill a process on your system that you

The /proc directory stores real time information about your PC, such as memory resources and CPU information.

can't otherwise end. The command to do this is (you guessed it): `kill`.

As an example we'll start a program from the menus and then kill it from a terminal. Both Red Hat and Mandrake come with KDE applications, so for this exercise start the KDE Advanced Editor. You'll find this in the Applications → Editors menu



for both Red Hat and Mandrake. Back in the terminal type:

```
ps x
```

Near the bottom of the list you should see an entry for 'kwrite', which is KDE's Advanced Editor. Now, suppose you don't want to run this program anymore, and suppose you're unable to close it via the File menu or by clicking on the close window button. In the process list you should see a category at the top called PID (Process ID). If this isn't visible, scroll the terminal up using the scroll bar. Every process on your system can be identified by an ID number that is assigned when the process starts. Under this column, find the number that corresponds to kwrite. Now for some fun; type:


```
kill [process id]
```

and it will be no more. If the process proves to be stubborn, you can force it to be killed outright by using `kill -9 [process]`. This rather impolite method doesn't allow the program to try and exit gracefully on its own first. Be careful with this command; be sure to kill only those processes you don't want. Don't try killing any processes that are essential to your system's well-being, or you will certainly run into trouble.

An alternative to the `ps` command is the program `top`. This program will display a similar output to that of `ps` but will update the list at a specified interval. It's handy if you want to monitor how programs behave. For example, it will keep an eye on how much memory and CPU time they absorb.

The `top`
program

BACKGROUNDING



You don't have to run the Advanced Editor from the menu. You can run it from a terminal; simply type `kwrite` and it will load. Notice, however, that the terminal you were working in no longer responds (you can't type anything in). This is because `kwrite` is running from it, and each terminal you use can only run one program at a time. Unless, that is, you tell Linux to run it independently. This is called backgrounding. If you want to run any program in the background, just append an amper-

sand (&) after the command.

First, press Ctrl-C to cancel `kwrite`. You should be able to use the terminal again. Then, type:

```
kwrite &
```

You will notice that you can continue using the terminal after `kwrite` starts. Type `ps aux` and you will see that `kwrite` is running as before.

There is another way you can do this. As noted above, you can use Ctrl-C to cancel a program that is running in the foreground (occupying the terminal you started it from). However, by pressing Ctrl-Z, you can specify that a program be backgrounded after it starts. This pauses the process and gives back control in the terminal. You can then issue the command `bg` to background the process. This has the same effect as originally running the program in the background.

QUICK HELP

You may already be wondering how you might go about finding the many switches that commands such as `cp`, `mv`, `rm`, `cat`, `ps`, `kill` and others accept. Linux has built-in help for all commands and programs on your system. In a terminal, type:

```
man ps
```

This will load the 'manual' for the `ps` command. As you can see, there are plenty of available switches. If you're ever in doubt about how a command should be used, just read its manual.

If you're the type of person who doesn't read manuals, now would be a good time to start. Just about everything you'll ever need to know about Linux is already on your system in the form of manual files and HOWTO guides. If you don't find the answers you need in this Pocketbook, you will find them in the installed documentation on your system. You can find more about help on page 164.

THE HOME DIRECTORY

Each user on your system is allocated a home directory, specified by the alias `~`. No matter which user you login as, typing `cd ~` will take you to your home directory. Home directories are stored under `/home` along with FTP and HTTP root directories, if you're running these services.

Just about everything you'll need to know about Linux is on your system in the form of manual files and HOWTO guides.

The home directory is not only a personal holding space for a user's files; it is also where Linux will store program settings related to the user. For example, each home directory will contain configuration files for the bash shell, newsreaders and email clients as well any other programs that decide to store user-specific information.

In future you will find yourself using the home directory to store personal data files and configuration files for the applications on your system.



THE HOSTNAME COMMAND

If you didn't set up networking during installation, when you reboot you may find your machine is called 'localhost.localdomain'. There is a good chance you'll want to rename your machine so as to give it a personality of its own. You can do this by typing:

```
hostname [your new hostname]
```

You will need to be logged in as root in order for the command to work. Alternatively, you can use the `linuxconf` program to set it. You will learn more about `linuxconf` later in the book.

EXPLORING THE FILESYSTEM

If you have performed a preliminary expedition into your Linux filesystem, you may have noticed that things are structured somewhat differently to other operating systems. Most new users will find the directory hierarchy confusing; however there is some logic behind the directory structure, and once you understand it, navigation becomes much easier.

Linux systems have a few directories that are almost always present. Below is a list of these directories and what is usually contained within them. You should note that these folders will often reappear many times throughout a system. Use the general concepts presented here and you should be able to work out what a specific folder is likely to contain. For example, `/bin` directories will always contain binaries (executable programs) while `/lib` directories will contain library files, and so on.

/

As mentioned earlier, the `/` directory (known as root) is the top-most directory on your Linux system. All other directories descend from this directory.

/bin

The `bin` directory is where Linux systems keep their executable files. Any user can run programs in a `bin` directory.

/boot

As you might have guessed, the `boot` directory is where information and files needed at boot time (such as the kernel) reside. On most Linux systems, the kernel is usually called `vmlinuz`, although `bzImage` is also common.

/dev

The `dev` (or device) directory is where Linux stores its device files. For example,

when a file is written to a floppy disk, it's written to `/dev/fd0`. There is also a special file called 'null', which is a bit bucket. Files written to `/dev/null` are immediately erased (not unlike a black hole).

/etc

The `etc` directory is normally used to keep configuration files. If you need to find a configuration file, this is a good place to start. By convention, configuration files usually end in 'conf' or 'rc'.

/home

The home directory is where you will find a user's home directory, and is used to store their personal preferences and files. It's also the first directory you will find yourself in when you logon.

/lib

The `lib` directory holds shared libraries (or links to libraries) used by various programs. The idea behind shared libraries is that there are routines needed by many different programs, and rather than have every program include that code, it is written once and stored in a library. Programs can then link to these libraries and use that code.

/lost+found

The `lost+found` directory is where Linux systems keep any files restored after a major system malfunction. If you ever lose important data, you should check

the `lost+found` directory before taking any drastic measures.

/mnt

The mount directory is where external filesystems (such as floppy drives and other partitions on your hard drive) are mounted. You can actually mount to a directory anywhere on your system. However, it's wise to be organised and mount all your filesystems through the `/mnt` directory.

/proc

The process directory contains 'virtual' files that represent various aspects of your system. You will find many files pertaining to system hardware and software statistics as well as directories relating to specific parts of your system. System information programs often read their data from the `/proc` directory, and you can view the same information yourself by using the `cat` command on the files contained here.

/root

The root directory is similar to the user directories in `/home`, except that it only contains the root user's (the system administrator) home directory.

/sbin

The `sbin` directory is similar to the `/bin` directory, except that it is reserved for system administration programs. You will need to be logged in as root to execute these programs.

/tmp

The temp directory is where programs store temporary files. For example, after installing Linux you are likely to find a file called `install.log`, which is a log of the installation proceedings. You can use the temp directory for your own temporary files.

/usr

The user directory is something of a miscellaneous directory. It got its name from being a 'user directory', where user-related programs and files were stored. It has since evolved into a secondary hierarchy, containing any read-only files that aren't required to be on the root partition. You will find most X-based programs located under the `/usr` directory.

/usr/local

Although the local directory is a descendant of the `/usr` directory, it is the directory where most users install new software. Although it's not necessary, it's probably a good idea to install any programs that you download into this directory to keep your system organised.

/var

The var directory is where files are stored that may change over time. By comparison, files under the `/usr` directory usually remain static, and are mostly read-only files. Some of the most notable files under the `/var` directories are the spools, which queue mail files and print jobs. You will also find a log directory, where you can view various logs of system activity.





Tools of the trade

Linux is based on Unix, and to get the most out of Linux it helps to learn a little more about basic Unix commands and how they are used.

This section is devoted to examining the capabilities of the common commands you'll find most useful in your exploration of Linux.

LISTING FILES

Earlier we introduced you to the `ls` command. However, there are a number of switches that you can use to format the output which `ls` produces.

You may have noticed that `ls` doesn't differentiate between files and folders. The `-l` switch tells `ls` to list the contents of a directory in long format, allowing you to see the permissions (covered in the next section). The `-a` command tells `ls` to list all files, no matter what type. Thus the most frequently used invocation of `ls` is:

```
ls -la [directory]
```

In fact, `ls -l` on its own has become so widely used that both Red Hat and Mandrake now use an alias, `ll`, that runs the command `ls -l` for you. That is, `ll` and `ls -l` are equivalent.

Some other useful switches to use with the `ls` command include the following:

-s

The `-s` switch lists the files in a directory and prefixes the size of the file to the listing. File sizes are listed in kilobytes rounded up.

-S

The `-S` switch tells `ls` to list files based on size.

-R

The `-R` switch will recursively list directories. This will display the contents of all directories below the one where the command is issued.

The `//` command

```
lanilla@Pocketbooks ~/ll
total 152
drwxr-xr-x  2 root  root    16K Oct 29 06:30 ll/
drwxr-xr-x  1 root  root    479 Oct 29 12:26 ll/ll/
drwxr-xr-x 11 root  root   173K Oct 29 12:27 ll/ll/ll/
drwxr-xr-x 53 root  root   409K Oct 29 12:26 ll/ll/ll/ll/
drwxr-xr-x  3 root  root    53 Oct 29 06:33 ll/ll/ll/ll/ll/
drwxr-xr-x  2 root  root   252K Oct 29 06:30 ll/ll/ll/ll/ll/ll/
drwxr-xr-x  2 root  root    16K Oct 29 12:26 ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x  8 root  root   162K Oct 29 12:26 ll/ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x  2 root  root    0 Oct 29 12:26 ll/ll/ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x  7 root  root    35K Aug 24 03:03 ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x 52 root  root    0 Oct 29 06:00 ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x 14 root  root    22K Oct 29 12:27 ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x  3 root  root   467K Oct 29 12:26 ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x  3 root  root   572K Oct 29 12:27 ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x 17 root  root   282K Oct 29 06:34 ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/
drwxr-xr-x 18 root  root    37K Oct 29 12:12 ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/ll/
lanilla@Pocketbooks ~/ll
```



SYMBOLIC LINKS

Occasionally, as you're browsing your system you may come across a directory or file that seems to point to another, such as:

```
lrwxrwxrwx 1 root root 10 Oct 1 02:58 tmp -> ../var/tmp
```

As you can see, the arrow implies that the tmp directory (in this example, found under the /usr directory) points to the /var/tmp directory. This is a symbolic link. Any files copied to /usr/tmp will in fact be copied to /var/tmp. Performing an ls on /usr/tmp will display the contents of /var/tmp.

Another way to determine a symbolic link is to look at the permissions. If the first character is an 'l' (as it is in the example above) the directory or file is a link. The syntax for creating a symbolic link is:

```
ln -s [original file or directory] [symbolic link]
```

Why use symbolic links? They reduce data redundancy by allowing you to point one directory or file to another. They can also be very handy when it comes to organising your data, as you will discover when you learn to compile your own kernel on page 120.

Of course these aren't all the switches available to ls; for a more complete list type man ls.

MOVING FILES

When it comes to actually copying files, these are the commands:

cp

The cp command works much like the COPY command under DOS. All the rules apply here that apply under DOS; you can copy files in the current directory or between directories. For example:

```
cp pocketbook.html /home/httpd/html
```

will copy the file pocketbook.html into the /home/httpd/html directory.

mv

Moving files is just as easy. The mv command moves a file to a new location. Essentially, mv is just a cp followed by an rm. For example:

```
mv pocketbook.html /home/httpd/html
```

will move the pocketbook.html file to the /home/httpd/html directory. However, mv can be used for more than just moving files. DOS users may remember a command called RENAME that allowed you to rename a file. There's no such command under Linux — why would there be? Renaming a file is the same as moving it from one filename to another filename, and that's exactly how you rename files under Linux:

```
mv pocketbook.html index.html
```

REMOVING FILES

You know how to list files and you know how to move them, but what about deleting them? The rm command will remove any file you specify. Be careful with this command, as rm has powerful switches that belie its seemingly simple nature.

-r

The recursive switch tells rm to remove files in all subdirectories where the command was issued. If you use a wildcard such as * then directories, as well as files, will be deleted. The subtle difference between rm * and rm -r * could have drastic results. Make sure you only issue this command when you want to remove a specific directory tree or collection of files. Luckily, Linux will prompt you to make sure you want to delete the files and folders you've specified — unless you decide to use the following switch.

-f

The -f switch will force any deletions, effectively bypassing safeguards. It is possible, therefore, to combine the switches to remove a directory hierarchy, and all the files therein, in one command:

```
rm -rf [directory]
```

Be careful with this command. If you use this on the root directory (/), logged in as root, you can wave goodbye to your system. You can even wipe the contents of mounted filesystems, such as your Windows drive, if it's attached at the time.

VIEWING FILES CAT AND MORE

Despite conjuring up images of cute

DIRECTORIES

Creating directories is achieved through the mkdir command. For example:

```
mkdir test
```

Similarly, you can remove directories using the rmdir command:

```
rmdir test
```

Alternatively, you can remove a directory and all its contents with the rm command (see 'Removing files' on this page). This is often the better method, as rmdir requires that a directory be empty before it can be removed.



THINK BEFORE YOU PRESS

Linux doesn't support undeletion. When you delete a file, it's gone for good. This is why safeguards such as prompts before deletion are wise — they are your only safeguards. Play it safe and work in a normal user account, not as root, and read what your system is telling you before you press that Y key.

felines, the `cat` command's full description is concatenation. Although it might seem strange to be viewing a file with a concatenation command, it's not so strange when you consider its original purpose was to concatenate multiple files for output to a terminal. View a file as follows:

```
cat [filename]
```

You can view multiple files by issuing multiple filenames. Like most commands, `cat` includes switches to alter the behaviour of the program. Two useful ones are `-n`, which will number the output from the command, and `-s`, which will suppress the output of multiple blank lines.

However, you may notice that when you view large files, they scroll by without giving you the chance to read them properly. There are other commands designed to allow you to better view files for reading, but in the meantime you can try:

```
cat [filename] lmore
```

The command at the end is actually two commands. The first, `l`, is a redirection command often called a pipe (you will learn about these later in this chapter). The second is the program `more`. By sending the contents of the file through the `more` program, you can read the file one screen at a time. To view the next screen of information, press the space bar. If you want to scroll line by line, use the Enter key. The 'B' key skips backwards and the 'Q' key quits.

HEAD AND TAIL

There will be times when you're only interested in peeking at a certain piece of information in a file, without viewing the entire file. For example, change to your home directory (`cd ~`) and type:

```
cat .Xdefaults
```

The `.Xdefaults` file lists a number of default settings to be used with X. Note

the capital X and the full stop preceding it. If you're getting used to using the Tab key, you probably just typed `cat .X [Tab]` and let Tab complete the filename. If the file is too long, and you're only interested in what the file might contain, you can use the `head` command to view the first 10 lines:

```
head .Xdefaults
```

You can also specify exactly how many lines you want to view by issuing the number as a switch:

```
head -20 [filename]
```

The `tail` command is just like the `head` command, except that it only lists the last 10 lines of a file. As with `head`, you can also specify the number of lines to display.

Keep in mind that most commands under Linux are highly flexible and have a range of switches you can use with them. Just check the man pages for any given command (`man [command]`) to find out more.

FINDING FILES

There are many ways to find files on your system, but the three simplest commands are as follows:

WHEREIS

This handy command will show you all file instances of a name you provide, if you're lucky. The syntax is:

```
whereis [filename]
```

Didn't we say Linux was easy?

LOCATE

The `locate` command will quickly provide the full path on your system for any file or directory that matches your search. The syntax is:

```
locate [filename]
```

It's wise to update the database `locate` uses regularly to ensure it can find new files on your system. The command that does this is:

```
locate -u -e /mnt
```

The `-u` switch updates the database and the `-e` switch tells `locate` to ignore directories under the `/mnt` directory — you don't necessarily want to index files on other filesystems, especially those that are only temporarily mounted!

WHICH DIRECTORY?

Getting lost in your filesystem? No problem, we have the command for you:

```
pwd
```

With print working directory you will always be able to find where you are. It's also possible to configure your prompt to display your full directory path, as well as many other values. If you're interested in tailoring a prompt, type `man bash` and read about the PS1 variable.

**DISK USAGE**

You may have noticed that the `ls` command doesn't automatically report disk usage like DOS' `DIR` command does. So how do you find out how much space is being used by files and how much is free? There are two commands you can use to determine disk usage on your system:

`du`

The `du` command reports your disk usage and displays the size (in 1K blocks) of each directory under the directory where the command was executed. If you run it from the root directory, you will see the size of every directory on your system, including any mounted filesystems (such as your Windows drive).

`df`

The `df` command reports disk free space and displays a summary of used and free disk space for every partition and filesystem mounted.

FIND

If you're looking for a specific file, you can use the `find` command. The syntax is:

```
find [directory] -name [filename]
```

If, for example, you were trying to find where `lilo.conf` file is stored, you can search from the root by typing:

```
find / -name lilo.conf
```

In this instance, the `find` command would also search any mounted drives, because they are part of the filesystem under the `/mnt` directory which branches off from root. This is handy if you want to search there, but will slow you down if you don't. In this case you can add the `-mount` switch to tell `find` to ignore mounted filesystems. You can, as always, type `man find` to learn more.

EDITING FILES

There are, as you would expect, many editors for Linux. In the original Linux Pocketbook we introduced you to `vi`, an old and reliable console-based editor. In these modern times, especially with the development of Gnome and KDE, there are plenty of graphical editors to choose from. That's not to say console editors

like `vi` don't have a place — if you're ever unable to enter into X you'll need to know how to edit files from the console. Long-time Linux users, including the Pocketbook team, still use and will continue to use editors such as `vi` as their standard editor, simply because it's quick and easy to use from the console (that is, if you've grown used to the command set).

Editing files, as you will discover, is core to using and making the most of Linux. Configurability and flexibility are the strength of Linux. Everything that can be tweaked, toggled, changed or adjusted in Linux is done through configuration files. In many cases there are graphical programs that help you configure Linux without messing with the configuration files that lie underneath, but more often than not you'll find yourself editing a configuration file in that most staple of operating system programs — the text editor.

Earlier we introduced you to the `kwrite` program, a powerful KDE text editor. If you find the need to edit a file while you read this book, simply start `kwrite` from the Gnome or KDE menu or run it from a console as follows:

```
kwrite &
```

Or, to start the editor and load a file to edit:

```
kwrite [filename] &
```

Given that the humble editor is

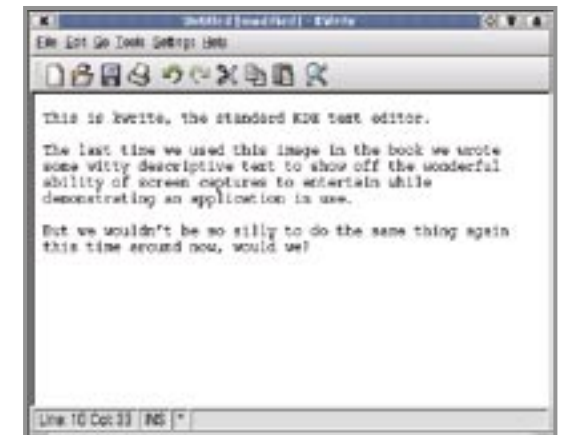
central to using the penguin operating system, Linux users are spoilt in their choice of editor.

Both Red Hat and Mandrake come bundled with a variety of editors to choose from. You will find the standard complement of console-based editors such as `vi` and `emacs` (simply run `vi` or `emacs` from a console), as well as graphical editors such as the KDE-based `kwrite` and `kedit`, or the Gnome-based `gedit` and `gnotepad` (run `gedit` and `gnp` respectively).

Explore the text editors that come with both distributions and find one that suits you.

If you need more powerful features, such as binary, coding or HTML editors you'll find plenty of these either already packaged with the distribution or included on the Linux Pocketbook CDs. If you're looking for a powerful word processor, check out OpenOffice, also found on the Pocketbook CDs.


The
kwrite editor





Permissions

Linux is a secure multiuser system that uses permissions to determine which users can access what resources. Sooner or later, every Linux user learns the joy of permissions.



When you list directories in long format (the `-l` switch), you'll be presented with a long list of information. What does it all mean?

The first column on the far left represents permissions. The first character is a marker to define what the object is. For example, `'-'` for a file, `'d'` for a directory, `'l'` for a link, and so on. The next nine characters define access permissions.

By default, a standard Linux user account doesn't have access to the entire system. Some directories and files are inaccessible, or even invisible, without the correct permissions. The

root user, as system administrator, is the only user who has full access to the entire system.

Permissions are broken down into three categories as seen in the table below.

Every user on a system belongs to a group. A group is simply a method of relating users and makes it easier for system administrators to manage things such as permissions for an entire set of users at once. Any permissions set on a file for a group automatically apply to all users belonging to that group.

The `ls -la` command reveals the permissions column on the left.

CATEGORIES OF PERMISSIONS

- Owner** Also known as the 'user', this is the person or process that created a file
- Group** The primary group that the user who created the file belongs to
- Others** Any other users who don't fit into the two categories above

For each category, one or all of the access permissions can be toggled for both files and directories.

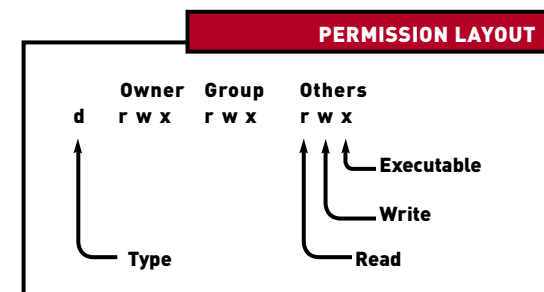
The diagram (right) shows the type of permissions and categories displayed when you list a directory in long format.

If you're wondering why you need to concern yourself with permissions, keep in mind that Linux is a secure multiuser platform designed to service many users. While you may not have a need to use these varied permissions yourself, you need to be aware of how they work. Later you'll discover how some services can only be operated from the root account (to ensure that sensitive services are accessed only by those qualified to have access).

You can change permissions the new, trendy way using your favourite GUI file manager (just right-click on a file and check out the permissions), or the traditional way using the command line. Just because we have such a strong sense of adventure, we'll show you how to do it the way real users do it.

CHMOD

You can change permissions (or modes) for an object on the command line using the `chmod` command. The



syntax is easy, and is based around adding, removing or directly setting specific permissions. Each permission category and value is represented as follows: `[u]ser`, `[g]roup`, `[o]thers`, `[a]ll`, `[r]ead`, `[w]rite`, `e[x]ecute`.

ADDING AND REMOVING PERMISSIONS

To add or remove permissions you can use the `+` and `-` switches to change the state of a permission. For example, to add executable permission to all categories of a file you would use the following command:

```
chmod a+x [file]
```

To remove write and executable permissions for the group and other users only, you would use:

ACCESS PERMISSIONS

PERMISSION	FILE	DIRECTORY
Read [r]	Read the contents of a file	List the contents of a directory
Write [w]	Modify the file	Add and remove files inside a directory



EXAMPLE PERMISSIONS

PERMISSION	ACCESS
<code>-rwx-----</code>	Only the owner can read, write and execute the file
<code>-rwxrwx---</code>	The owner, and anyone in the group the owner belongs to, can read, write and execute the file
<code>-rwxrwxrwx</code>	Everyone can read, write and execute the file, including the owner and anyone in the group the owner belongs to
<code>-rwxrw-r--</code>	The owner can read, write and execute the file; users in the owner's group can read and write, but not execute the file; everyone else can read the file, but not write to it or execute it
<code>drwxr-xr--</code>	The owner can list, manipulate and search in the directory; users in the owners group can list and search in the directory; everyone else can only list the contents of the directory

```
chmod go-wx [file]
```

It's important to note that adding or removing permissions this way will only modify the current permissions. For example, the above command would not remove writeable and executable permission from the owner if it were present.

By contrast, setting specific permissions allows you to set permissions by category (keep up, this does make sense!)

SETTING SPECIFIC PERMISSIONS

To change, for example, the permissions on a file to set the user (owner) and group categories to have only executable permission, and leaving permissions for other users alone, you would use:

```
chmod ug=x [file]
```

To change permissions recursively for a directory (and thus all the files and directories within it) so that:

- 🔑 the owner has read, write and execute permission;
- 🔑 the group has read and execute permission;
- 🔑 other users don't have any permission to access the object at all; you would use the command:

```
chmod -R u=rwx,g=rx,o= [directory]
```

It goes without saying that you can only change permissions of files and directories that you have permission to change. The root account can change permissions for any file or directory on the system.

DEFAULT PERMISSIONS IN RED HAT AND MANDRAKE LINUX

Red Hat and Mandrake use a system called User Private Groups, which is a way of ensuring that files created by each user are private by default. When a new user is created, a new group is created with the same name as the user, with that user as the only member.

Sound confusing? If we added a new user called 'pocketbooks', for example, a group would also be created called 'pocketbooks'. The only member of the group then would be, you guessed it, 'pocketbooks'.

In Linux, each user has a set of default permissions (known tech-

It goes without saying that you can only change permissions of files and directories that you have permission to change.

nically as a umask) for files they create. For Red Hat and Mandrake Linux, this umask is set to allow full access to the user and the user's group, and no access to other users.

What has been covered in this section is enough for day-to-day tasks, but there is far more to permissions than there is space here to cover. The Advanced Linux Pocketbook will delve deeper into permissions and show you how to change more than just the access permissions of files and directories. In the meantime, you can learn more with `man chmod`.

CHANGING USER AND GROUP OWNERSHIP

You'll notice when you list a directory in long format that the name of a file's owner and group are displayed. These, too, can be changed if required, but only if you're logged in as root. For this you use the `chown` command. For example:

```
chown -R pocketbooks.linuxpb /usr/share/data
```

would change `/usr/share/data` and all files and directories within to be owned by the user 'pocketbooks' and the group 'linuxpb'.



Shell operation

Now that you know the basic commands and how to use them, it's time to consider how the shell operates and what rules it follows.

WILDCARDS

As with most operating systems, wildcards can be used to represent files and directories. Those of you who remember DOS will recognise these wildcards, but be aware that they operate slightly differently under Linux.

THE * WILDCARD

Probably the most useful wildcard is the asterisk. It represents everything, no matter how many characters might make up a filename. For example, copying the entire contents of one directory to another:

```
cp old/* new
```

Even if you want to narrow down the files you want to work with, the asterisk can still be helpful. For example:

```
cp old/a*.conf new
```

This command would copy all files from the old directory that begin with 'a' and have the extension 'conf' to the new directory. Under DOS, the same could be achieved by using '*.*'. Under Linux however, '*.*' will only match those files that have one or more extensions; files without extensions won't be picked up.

Although it is a very handy wildcard, be careful with it. Using it in an `rm` statement will tell `rm` to delete everything.

THE ? WILDCARD

The question mark wildcard works the same as the asterisk, except that it will only match one character. For example, if you wanted to copy just the C program files, header file and object files for a program called `hello` to a 'hellodev' directory, you could use the following:

```
cp hello.? hellodev
```

This would match `hello.c`, `hello.o` and `hello.h`, but not `hello.bak`.

THE [...] WILDCARD

You can also define a specific subset of characters to match. For instance, to list only the files that begin with 'a' or 'A', you could use the following :

```
ls [aA]*
```

This presents another way to copy program files. For example, to copy all program files that end in `.c`, `.o` or `.h` to a directory called `C`, you could use the following command:

```
cp *.c[h] C
```

Finally, you can use the brackets wildcard to remove multiple consecutive

characters without typing them all in individually. For example, to move all of the files that begin with the letters 'a' through 'm' to a directory called 'temp', you could use the following syntax:

```
mv [a-m]* temp
```

All of these wildcards will match directories as well, so be cautious when using them.

DIRECTORY SHORTCUTS

If you have used a command line before, such as in DOS, you will know that there are various constructs the filesystem abides by. One of these is the current directory. Initially, it can be somewhat confusing to grasp the concept of the current directory if your only experience has been graphical interfaces.

RUNNING PROGRAMS FROM THE CURRENT DIRECTORY

Whenever you issue a command, the `PATH` environment variable under Linux specifies the directories that should be searched, in corresponding order, to determine where it might be found.

That's how you can run commands such as `ls` anywhere on your system, despite the command residing in the `/bin` directory.

Linux is often configured so that the current directory (`.`) is not included in the `PATH`. This means that unless you specify the full path to the current directory, you can't run a program from there. There's a very good reason for this — by restricting the commands you run to those in the `PATH`, you can't run a potentially harmful program by accident.

Let's say you just installed a new version of `loadmeter`, a nifty CPU monitor program. You might have installed it to `/usr/local/bin` and despite changing to the `/usr/local/bin` directory, you discover that typing `loadmeter` doesn't start the program. All you need to do is prefix it with the current directory:

```
./loadmeter
```



COMMAND COMPOUNDING

You can execute multiple commands on one line by separating each with a semicolon (;). One command will start as soon as the command that precedes it finishes, regardless of the results of the command. For example:

```
cd /; mkdir test; ls -la
```

You can stack as many commands as you want this way. Alternatively, using a double ampersand operator (&&) between each command will only run each command in order as long as the one preceding it completes successfully.

The '.' directory

Every directory in Linux contains at least two subdirectories: '.' and '..'. The dot directory (.) acts as a synonym for the current directory. For example, the following two commands are identical:

```
ls
```

```
ls .
```

The '..' directory

The '..' directory is a link to the parent directory of the current directory. For example, assuming that we're in the /usr/local directory, the '..' folder refers to the /usr directory.

Because these directories exist in every directory on your system, you can use them as a shortcut in your commands. For example, if you were in the /usr/local/tmp directory and you wanted to copy all the files from /usr/local/tmp to /usr/local/bin, you could use the command:

```
cp * /usr/local/bin
```

Or, you could use:

```
cp * ../bin
```

PIPING AND REDIRECTION

A pipe is, as the name implies, a connection through which data can flow. When it comes to the shell, a pipe is used to connect two programs.

The | operator

The | (this is the key next to the square right bracket on most keyboards) operator represents a pipe. This links the output of one program to the input of another. For example:

```
cat README | more
```

This would cat the file README and pipe its contents through the more program, which would allow you to view it page by page.

Redirection with the > operator

Redirection is the process of taking something that's headed for one place, and redirecting it to another. For example, the ls command lists the files in the current directory and prints to the screen. Instead of printing output to the screen, however, we could output to a file. This is accomplished by using the greater-than sign (>).

To output the directory listing of your root partition to a file, you could issue the following:

```
ls / > rootfiles
```

Note that nothing would be printed to the screen. However, typing ls would reveal a new file called 'rootfiles'. If you then cat this file you could view its contents.

Redirection with the >> operator

Sometimes you might want to append the output of a pipe to a file instead of overwriting it. This can be

accomplished using the >> redirection operator. If you issue the command:

```
ls -r >> rootfiles
```

you would notice that the new listing has been added to the end of the file 'rootfiles', rather than replacing it.

Redirection with the < operator

Not only can you redirect file output, but you can also redirect file input. When programs rely upon keyboard input for their operation, you can substitute a file containing information using the < operator.

For example, you might have a script to clean up your work directories (called, for instance, 'cleanup'). In order to respond 'y' to all the delete confirmations, you create a file called 'answer' that is filled with nothing but 'y' on each line. You then use this as input for your script, so you don't have to manually respond:

```
cleanup < answer
```

This idea — having to answer 'y' to programs — is common, and fortunately Linux comes with a nifty program that does just that. It's called yes. Another way of passing the required 'y' answer to your script would be:

```
yes | cleanup
```

The yes program produces a continuous stream of 'y' letters.



Mounting filesystems

If you have multiple operating systems on your PC, or you want to access removable media and network resources, you'll need to learn the art of mounting.

As discussed earlier in this Pocketbook, mounting is the process of attaching external filesystems to the Linux filesystem hierarchy. For example, attaching a CD-ROM or Windows partition to become an extension of the Linux filesystem tree, thereby creating a seamless directory system that spans all local and remote resources.

Although mounting is mostly automated in today's Linux distributions, it is, like many of the topics covered in this chapter, integral to using Linux and essential to understand if you're interested in mastering Linux.

ALL ABOUT MOUNTING

To mount a filesystem you use the mount command:

```
mount -t [filesystem] [device] [mount point]
```

The '-t' switch specifies the filesystem type that should be used, such as VFAT (for Windows FAT16/32 partitions), Ext2 (for Linux), ISO9660 (for CD-ROMs), and so on. The 'device' switch specifies the device you want to mount, which for local devices will exist under your /dev directory. Lastly, 'mount point' specifies where you want to attach the device on the Linux filesystem tree.

Before you can access your CD-ROM, for example, you will need to mount it. By default, most Linux distributions come with an entry under the /mnt directory called 'cdrom'. You can mount your CD-ROM like so:

```
mount -t iso9660 /dev/cdrom /mnt/cdrom
```

You can then list, copy and work with the contents of the CD-ROM through the /mnt/cdrom directory.

If you want to mount a filesystem which doesn't have an entry in /mnt, you can simply create one. In fact, you can mount a filesystem to any empty directory on your filesystem. However, it's wise to mount all your filesystems under /mnt as it allows you to organise your data and keep track of external filesystems. It's

also handy when it comes to using commands that operate on or exclude your mounted filesystems.

Assuming your Windows C: drive were on device /dev/hda1, you could mount your Windows partition thus:

```
cd /mnt
mkdir cdrive
mount -t vfat /dev/hda1 /mnt/cdrive
```

You would then be able to access and manipulate the contents of C: as easily as any other directory on your system. Take note that you need to be in root in order to mount devices, so be sure to su first.

Mounting isn't limited to just local resources such as floppies, CD-ROMs, partitions and extra hard drives. You can also mount network resources such as SMB (Windows) shares and NFS (Linux) shares.

THE EASY WAY

Both Red Hat and Mandrake Linux will automatically mount any CD-ROM or floppy you insert. Red Hat will also set up /mnt entries and automatically mount some removable media (such as Jaz drives), while Mandrake will do the same for any other partitions (such as your Windows partition) it finds on your hard drive. These will be mounted as /mnt/win_c, /mnt/win_d, and so on.

In addition, both distributions will create desktop icons which allow you to easily access and manage these devices. These simply link the file manager to directories under

/mnt. You can, for example, use a file manager to browse /mnt to access mounted filesystems.

THE MAGIC FILE: FSTAB

Just when you thought you'd managed to escape cryptic Linuxisms in this section, along comes the fstab file.

Essentially, the fstab file contains information used by mount to help you mount devices, as well as define devices to be automatically mounted — and even checked — at boot time. Type cat /etc/fstab to see what the file contains. You'll notice a number of entries that define your root partition, swap and standard removable devices, such as the CD-ROM.

The format of each entry lists, in order: the device, mount point, filesystem, options, dump order, and fsckorder.

The first three are the same as supplied to the mount command. 'options' allows you to specify special options for the mount. For example, CD-ROMs are mounted with the read-only option, and the root partition is specified to automatically be mounted at boot. 'dump order' tells Linux backup utilities that you'd like the filesystem backed up. 'fsckorder' tells Linux if and when you'd like the file system checked if the machine isn't shut down properly. Linux can check and repair errors in a wide variety of filesystems, including Windows partitions. Your root partition should always be set to 1.

Any devices specified in the fstab file which aren't mounted at boot



time, such as the CD-ROM, can be mounted easily from the `/mnt` directory simply by typing `mount [mount point]`, thus bypassing the need to specify filesystem type and device. This same information is read by Gnome and KDE when mounting removable media.

You can, of course, add your own entries to the `fstab` file. For instance, under Red Hat, any Windows partitions you have on your system aren't immediately accessible. To enable them to be mounted at boot time like Mandrake you could add the following:

```
/dev/hda1        /mnt/win_c vfat defaults 0 0
```

This mounts the first partition of the first IDE drive as a VFAT (FAT16/32) filesystem to the `/mnt/win_c` directory. To see all partitions available on your system for mounting, type `cat /proc/partitions`. After saving the file, create a directory called 'win_c' under `/mnt`; and the next time you boot up, `/mnt/win_c` will be available to use.

Be careful when permanently mounting a filesystem this way, because any system-wide commands you run from the root (`/`) directory, such as a search or a deletion, will also apply to all partitions and devices mounted under `/mnt`. For example, if you run a recursive deletion command from the root directory, it will recursively delete across the Linux partition and any mounted partitions.

UNMOUNTING

Removable media, such as floppies and CD-ROMs, need to be unmounted before you can remove them from your system. There's a good reason for this: Linux optimises disk I/O by buffering information, and this includes writing


data to removable media mounted on its filesystem hierarchy. Unmounting merely insures that all information that needs to be written has been written and that the device is detached from the filesystem tree.

Gnome and KDE will automatically unmount a device if you eject it, but you can also manually unmount from the command line by typing:

```
umount [device]
```

As always, for more information about mounting, type `man mount`.

The `/mnt` directory.



Installing new applications

It's only natural that you'll want to plunge into the depths of the Web in search of new applications. But once you've got them, how do you install them?

FILE FORMATS

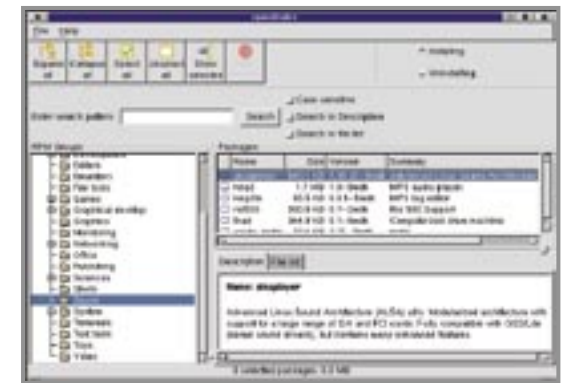
Programs under Linux come in a variety of different formats. Some are in a standard package format — the same as the packages you installed for Linux — while others are in a compressed format called tarballs, not unlike Zip for DOS and Windows.

Because Linux applications are more often than not distributed as open source, you can download the source code and compile the program yourself. Sometimes this has advantages — for example, compiling programs to take advantage of your Pentium CPU by using Pentium optimisations.

Many advanced Linux users prefer to 'cook for themselves' this way and install their applications from source. But installing the ready-to-go, precompiled versions (called binaries) is much easier. In the name of Linux wisdom, we'll show you how to do both.

PACKAGES

There are all kinds of package formats for different distributions (for example, the DEB format for Debian), but the format used by both Red Hat and



Mandrake is called RPM, or the Red Hat Package Manager format. RPM packages not only contain the compressed application (as with Zip files) but also information on how the program should be installed, who made it and where it came from, the application's dependencies, and much more.

RPM is more than just a package format, however — for every RPM file you install and remove, a centralised database on your system is updated to reflect the change. At any time you can query the database for a full list of installed applications (see the 'Command reference' on page 172), and the `rpm` program will use the

Mandrake's RPM manager is a well-designed front end for `rpm`.



database to ensure all your installed applications work well together and don't accidentally break one another when they are installed or removed. This prevents strange application behaviour, such as programs no longer working, the likes of which is often experienced under Windows.

Packages can be installed using a Gnome or KDE application, or from the console. For Gnome and KDE you'll find that the built-in file manager will either start an RPM manager for you, or at least let you view inside an RPM when you click on an RPM file. When it comes to installing and uninstalling RPM files you can use `gnorpm` for Gnome and `kpackage` for KDE. If you're running Mandrake, the `rpmrake` utility is an excellent alternative for either of the above programs.

To install, upgrade or remove a program from the command line simply open up a new terminal, log in as root using `su`, and then run:

```
rpm -ivh [package]    To install a package
rpm -Uvh [package]   To upgrade a package
rpm -e [package]     To remove a package
```

What are the switches? 'i' and 'U' represent 'install' and 'upgrade' respectively; 'e' is to 'erase'; 'v' produces a verbose output; and 'h' prints hash symbols (#) as a form of progress meter (so you know `rpm` is doing something).

Sometimes when you try to install a new program the RPM manager or `rpm` itself may report that it was unable to install the program due to failed dependencies. This means that either a) the program you are installing requires another program to operate, and this program needs to be installed first, or b) a program currently installed relies upon an older version of the program you are trying to replace. This is `rpm` doing its job and stopping you from breaking currently installed applications at the expense of installing your new one.

The solution to both of these is to install or upgrade the missing dependency before installing the new program. If you don't already have the packages you need, you can surf to `Rpmfind.Net` (www.rpmfind.net) and search for the program name `rpm` is reporting that it needs installed first. If you need a number of packages to be installed, and you have them all downloaded, you can let `rpm` work out the correct order to install the programs by supplying `rpm` with all the filenames on the one command line. Alternatively, simply select all the downloaded packages in an RPM manager and it should do the same.

TARBALLS

Tarballs are so named because of the type of archive programs used to create them. Whereas Windows has formats such as Zip, Lha, Arj, and so on (all of which you can also use under Linux!), the core format for Linux compressed archives is the tarball, represented by files ending in `.tar.gz`, or `.tgz`. The

reason there are two extensions to the compressed file — `.tar` and `.gz` — is that two programs are used to create them.

Linux uses a form of Zip known as `gzip`, or GNU Zip, to compress files. GNU Zip works exactly the same way as Zip under Windows and creates files with a `.gz` extension. Under Linux an extra step, known as tarring, is often performed before zipping files. A program known as `tar` (the Tape ARchiver, originally designed to store files on tape) first concatenates files into a single volume ready to be compressed by Zip. Doing this can dramatically increase the efficiency of the Zip algorithms and thereby create even smaller archives than zipping alone. The resulting files are affectionately called tarballs.

You can decompress tarballs using a file manager under Linux by just clicking on the archive to open it and then work with the files as you would any other directory. Alternatively, you can decompress tarballs from the command line like so:

```
tar -zxvf [archive]
```

The switches, in order, represent 'z' to first unzip the archive; 'x' to extract (or 'untar'); 'v' to be verbose in what `tar` is doing; and 'f' to specify the archive filename.

It's always a good idea to download or copy tarballs into your `/tmp` directory to work with. This way, whether they are binaries or source, you can decom-



press them and install or compile the application from your `/tmp` directory, deleting the directories created by the archive when you have finished.

Whether you're installing source or binaries from tarballs you should always immediately read any README and INSTALL files located within. Whereas installing an RPM will automatically copy the program's files to all the correct locations in your filesystem, decompressing a tarball simply decompresses it. You still need to install the program by following the instructions in the README and INSTALL files. You'll notice these files are almost always capitalised — this is to draw your attention to them. The README file will often detail the program, how to use it and, sometimes, how to install it. More often than not you'll be directed to read the INSTALL file. Make sure you always read these files to ensure you've installed and configured the application correctly.

CUSTOMISED INSTALLERS

Some applications (such as OpenOffice, and a few of the Loki games) use Windows-style GUI installers. While

You can use KDE's Konqueror to explore tarballs as if they were ordinary folders, using drag and drop to extract their contents.



these programs are self-installing, sometimes you may find that running the installer name doesn't start the program. This happens because Linux hasn't been told you have permission to run it (if it was copied across from a source, such as a CD where Linux permissions might not be present, for example). If this happens simply run `su` to login as root and then issue the following command: `chmod a+x [filename]`. This marks the file as executable for all users, and you can then run the program (as root or as your user account) to install it.

COMPILING SOURCE

Source code can come in both file formats — as RPMs (known as Source RPMs, or SRPMs) and, of course, as tarballs.

Source RPMs have, interestingly, the ability to let you compile the source code within them and then build a binary RPM file ready to be installed. You can do this as follows:

```
rpm --rebuild [filename.srpm]
```

This will create a binary RPM in the `/usr/src/RPM/RPMS` directory under the architecture for your machine (i386, i486 or i586). You can then change to this directory and install the RPM with `rpm -ivh` as discussed earlier.

If you want to do things the hard-core geek way and compile the source manually, here's how it's done (it's actually very easy).

After decompressing the tarball from your `/tmp` directory, change to the directory that's created and first — as mentioned previously — read the README and INSTALL files. They might suggest you edit the certain configuration files to personalise the source code to your system. Do so if you need to and then, when you're ready, compile the source like so:

```
./configure
make
```

There's no real standard for compiling source. It all depends on how helpful the authors have been in automating the process with a configure script. If the program compiles well, you can then install it with:

```
make install
```

You will, of course, find a fine selection of applications and other cool stuff on the Pocketbook CDs, some of which are in RPM format and some of which are in the classic tarball format. Explore and enjoy!

Compiling the kernel

Ever wanted to build your own operating system? Now you can. Learn how to tailor Linux for your machine and reap the performance benefits.

This is one of the most impressive and exciting aspects of Linux — building your own kernel. All operating systems come with a kernel; it's the very core of the OS and it handles the management of your hardware and software. The difference between Linux and commercial operating systems such as Windows is that Linux gives you the source code — as with all open source programs — to the kernel itself. This enables you to build your own kernel and configure it for your specific system.


By default, Windows and Linux will ship kernels designed to support the widest range of hardware and features out of the box. While such compatibility allows the operating system to support the biggest possible audience, it also means the operating system is designed to support the lowest common denominator. Why run a kernel designed for a thousand different machine configurations when you can run one that's designed for your specific system? By creating a kernel tailored to your machine you can increase performance and save system resources.

With Windows you can't do anything about this, because you don't have access to the source code. With Linux, however, you can compile your own kernel tailored to your system.

Once upon a time it was necessary to compile your own kernel in order to maximise performance on a system, but current distributions often come with pre-compiled kernels optimised for a variety of processors, including Pentium CPUs. Red Hat 8.0 and Mandrake 9.0 both do this, and an optimised kernel will already have been installed on your system.

However, if you want to squeeze even more performance out of your system, or you just want to explore the kernel, then 'rolling your own' is the way to go.

Compiling your own kernel is not only a fantastic insight into the inner workings of Linux, it's also a lot of fun (if you're the geeky type) and easy to do. In this section we'll show you how to play with the kernel source that comes with Red Hat and Mandrake Linux as well as how to go about upgrading the kernel when newer versions are released.





1 STEP ONE

CONFIGURING THE KERNEL

If you followed our installation guide the source to the kernel will already be installed on your system. Open up a terminal, su as root, surf on over to /usr/src/linux, and type:

```
make xconfig
```

You will be presented by a graphical kernel configuration program that lists the different aspects of the Linux kernel (If you ever find yourself without X you can, alternatively, configure the kernel through the console by running make menuconfig.)

From here you can explore the vast range of features the kernel supports and decide which features you want to include or remove. Start at the top and cycle through every option of every topic. Pick and choose the features you want, and click Help on everything, even the options you think you understand. If you have any doubts about a feature, leave it on its default setting, or follow the recommendation Help provides. It's easy to build a new kernel, but there's no guarantee it's going to boot if you leave out a key ingredient.

If this is your first foray into exploring the kernel be sure to grab a coffee, as it's going to take you half an hour or more to run through all the configuration options. If you don't feel comfortable with compiling your own kernel, don't let that stop you exploring the features it offers. Until you actually replace the kernel after you've configured and then compiled it, nothing will change on your system.

You may notice in your explorations that many features sport an 'm' option. Much of the kernel is modularised — that is, some features can be loaded into memory and unloaded in real time as required. Any features flagged 'm' will be compiled as a module for this purpose. By contrast, anything marked 'y' will

Configuring the kernel with xconfig.



be a permanent feature compiled into the kernel. Some parts of the kernel can't be modularised because they are critical to its operation, while others might only be available as a module alone.

When you're done picking and choosing the hardware support and features for your system, simply click 'Save and Exit'.

2 STEP TWO

COMPILING THE KERNEL

Before compiling the kernel you first need to set up any necessary dependencies and clean up the source tree from any previous kernel compilations. To do so type:

```
make dep && make clean
```

If you're not sure what the double ampersands do, take a quick look at page 108. When these commands are complete you can compile the new kernel using the following command:

```
make bzImage
```

Be sure to capitalise the 'I' in there. Compiling the kernel can take a while, so top up that coffee while you wait.

When it's finished compiling there are two more commands you need to run:

```
make modules
make modules_install
```

The first compiles all the features you specified to be included as modules, while the second installs them for you under the /lib/modules directory (this is where the kernel loads them from). Depending on the configuration of your kernel, this too may take quite some time to compile. Take note that modules are not only loaded on bootup as required to support the hardware in your system, but can also be loaded and unloaded at any time while using

Linux, either by the kernel itself as it deems necessary or by you at the command line. Modules are a powerful operating model and the art of using them is covered in more detail in The Advanced Linux Pocketbook. With the kernel compiled, your next step is to install it. You'll find your new kernel conveniently placed in the /usr/src/linux/arch/i386/boot directory. Type ll arch/i386/boot and you'll see bzImage waiting patiently to be installed. To install the new kernel all you need to do is copy it to the /boot directory along with the necessary System.map file which was generated (in the /usr/src/linux directory) when the kernel was compiled. Copy them both as follows:

```
cp arch/i386/boot/bzImage /boot
cp System.map /boot
```

Lastly, you need to update the module dependencies tree, a small file that tells the kernel which modules depend on other modules to work. Simply type:

```
depmod -a
```

All that remains to be done is to update the boot loader to include the new kernel. Even if you're using a different loader like GRUB (Mandrake defaults to this), we'll cover LILO here since it's installed automatically with both distributions.

Head on over to /etc and edit the file lilo.conf using your favourite text editor. All we need to do here is add



the new kernel as an entry for LILO to boot. It's a good idea to keep your current working kernel as the default should your new kernel have problems. To add a new entry simply duplicate the default entry and then change the `label=` and `image=` lines to match your new kernel. In the example below we've inserted the new entry (italicised for your convenience) after the 'failsafe' kernel, but before the entry that boots the Windows partition:

```
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
vga=normal
keytable=/boot/us.klt
lba32
prompt
default=linux
timeout=50
message=/boot/message
image=/boot/vmlinuz
    label=linux
    root=/dev/hda2
    read-only
image=/boot/vmlinuz
    label=failsafe
    root=/dev/hda2
    read-only
image=/boot/bzImage
    label=new kernel
    root=/dev/hda2
    read-only
other=/dev/hda1
    label=Windows
    table=/dev/hda
```

The exact configuration of `lilo.conf` will differ slightly depending on your system and chosen distribution, but the principles are the same. Save the file and then register your new kernel by typing: `lilo`

LILO will display the entries it is registering along with a '*' next to the default boot option. If it reports an error chances are you've mistyped a line, so open the file and fix up any problems.

If you're using GRUB as a boot loader you'll need to edit the `menu.lst` file in `/boot/grub`, and then run `sh install.sh` to initialise the changes. Alternatively,

since LILO is installed with Linux by default, you can use LILO as a boot manager over GRUB by simply running it.

When you're ready shut down and reboot. If everything went well you will be able to select your new kernel from the LILO boot menu. If everything went really well you'll be able to successfully boot your new kernel as well. If after a while you're satisfied that your new kernel is working OK, you can make it the default by editing `lilo.conf` and setting the `default=` line to the label you gave your kernel and then run `lilo` again.

Just as with commands, you can get help on `lilo.conf` by typing `man lilo.conf`

GETTING NEW SOURCE

If you're one for keeping up with the times you can download and install new kernel versions as they are released on the Internet. The best place to do this, apart from keeping up to date with the news on Slashdot and

Freshmeat.net (see page 166 for more information), is the Australian Linux Kernel Archive, an Australian mirror of the central kernel repository. Using your favourite FTP client, just connect to `ftp.au.kernel.org` and check out the latest kernels from the `pub/linux/kernel` directory.

When it comes to choosing a kernel it helps to understand a little about the kernel naming convention. There are two types of kernel available: stable and development. Stable kernels are for everyday use and are shipped with distributions such as Red Hat and Mandrake. Development kernels are, as their name suggests, constantly under development and should not be relied upon for everyday use. Development kernels are the cutting edge of Linux and can be fun to play around with, as they contain the latest features and performance enhancements that will eventually make it to the next stable kernel release. You can differentiate kernel source by the name of the file archive (see below).

KERNEL NAMING CONVENTION

Linux kernel source is identified by major and minor revision numbers. For example, the stable kernel 2.4.19, is the 19th update to the 2.4 stable kernel.

The naming convention looks like this:

`linux-2.4.19.tar.gz`

↑ Minor version number of the kernel
 ↑ Even represents stable kernels, odd represents development
 ↑ The major version number
 — So you know this is Linux kernel source

At the time of writing, the 2.5 kernel was nearing completion. When it arrives, a new development tree for 2.6 will begin.

INSTALLING NEW SOURCE

Kernel source archives are typically quite large — near the 20M mark — and expand to be even larger source trees (over 100M in size), so make sure you have the space available.

The following guide assumes you've downloaded a new stable kernel, let's say version 2.4.20, to your /tmp directory. In order to hang on to the old source (in case you still need it), and because the kernel source tree always begins with 'linux', you'll need to do a little symbolic linking. Change to /usr/src/linux and ls -la; you'll see a directory containing the current kernel source and a symbolic link, 'linux', pointing to it. Make a new directory for your source, delete the old 'linux' link and then create a new link which points to this directory:

```
mkdir linux-2.4.20
rm -f linux
ln -s linux-2.4.20 linux
```

You should end up with something like this in your /usr/src directory:

```
linux —> linux-2.4.20
linux-2.4.20
linux-2.4.19
```

Now uncompress the kernel source from the /usr/src directory as follows:

```
tar -zxvf /tmp/linux-2.4.20.tar.gz
```

Your new kernel is then installed! You can change to the 'linux' directory and configure and compile a new kernel. Take note that when changing to new kernel versions you should always compile new modules as well.

Configuring and compiling your own kernel is a huge learning experience that is well worth the effort. Take the time to explore all the features your Linux kernel offers, and if you're inclined to programming, browse the source code. If you think you can make a contribution to the kernel, write some code and submit it for inclusion. If it's good enough, the whole world could be using it before you know it!



Surviving in Linux

There comes a time for every Linux user where the following wisdom becomes invaluable. Bookmark this section, you'll be needing it.

SAVING SPACE

After you've been using your Linux system for a while you might find that you're in need of space. Freeing up space is a matter of deleting temporary files and applications you no longer need. The following command helps you find which directories are using space (and thus, the files within them):

```
du -h -x --max-depth=1
```

As you can see, this is a form of the `du` command you learned earlier. This command will show the size of directories from the level you run the command. If you run it from `/`, you will see the size of all the directories branching from `/`. If you run it from `/usr`, you'll see the size of directories under `/usr`. This way you can hunt down the culprit files that are taking up valuable space. As a guide, here are some places and files to look for.

Core files

If a program aborts, or you forcibly close X, a core dump will be produced. Search for these files and remove them; they're only of use to debuggers.

`/tmp`

This is the main directory where you'll find files that are taking up space. If they're not in use, delete them.

`~/.gimp`

If you've been using The Gimp frequently, you may find plenty of temporary (and often big) work files no longer in use.

`~/.netscape/cache`

Even if you use the 'Clear disk cache' option in Netscape there'll still be temporary cache files you can remove.

Whatever you do, always be sure to check that you are deleting a temporary file that's no longer needed. If you're not sure, don't delete it!

When it comes to removing applications you can use a package manager such as `kpackage` or the excellent `rpm-drake` from Mandrake to remove applications you don't need anymore.

RESCUE DISKS

If you didn't make a rescue disk during installation, now is the time to do so. With a rescue disk handy you should always be able to get out of trouble.



If your motherboard is capable of booting from a CD-ROM drive, all you'll need is the Red Hat or Mandrake CDs, as this allows you to boot a rescue disk from the installation screen. If you can't boot CDs you will need to make a rescue disk. Find the file `rescue.img` on the Red Hat or Mandrake installation CD and write it to a blank floppy using:

```
cat rescue.img > /dev/fd0
```

A rescue disk is a small Linux system that runs from a single floppy. It contains a minimal Linux kernel and the tools you need to get your system going again. In particular, it contains `lilo` to allow you to write a new `/etc/lilo.conf` and `fdisk` to allow you to look at the partition table. To use the rescue disk, reboot using the floppy or CD and type `rescue` at the prompt.

If the problem is that your LILO configuration isn't currently working and you know the location of your root Linux partition, you can boot it as your root partition and fix your LILO problems easily. If your root partition were on `hda3`, for example, you would type the following at the LILO prompt:

```
linux root=/dev/hda3
```

This means Linux will boot using the kernel on the rescue disk but will mount the root partition on `/dev/hda3`.

SINGLE USER MODE

Single user mode is a specific runlevel where only a few of the usual processes are started up. As the name implies, only a single user can use Linux in single user mode. It is particularly useful when one of the daemons started at boot time fails to exit properly and the boot hangs. To enter single user mode, type the following at the LILO prompt:

```
linux single
```

This assumes you're booting the entry 'linux'. If you've installed a new kernel under a different name you'll need to use that name instead. Once Linux has booted, you will be dropped into a shell as `root` and the `#` prompt will be displayed. Do whatever you need to do to fix the problem (if you're confident enough) and type `exit` to continue with the boot process.

Another special mode, which is useful when you are having problems with a filesystem, is emergency mode. It boots without mounting extra filesystems and the root partition is mounted as a read-only filesystem. To enter emergency mode, type the following at the LILO prompt:

linux emergency

REINSTALLING LILO

Sometimes your boot sector might be overwritten. This usually occurs when installing Windows, which blatantly overwrites the boot sector. If this happens you'll need to reinstall LILO into the boot sector by running `/sbin/lilo`.

If you have booted from a rescue disk, you'll first need to mount your root filesystem before installing LILO to the boot sector. If your root filesystem were on `hda3`, you would type:

```
mount -t ext2 /dev/hda3 /mnt
```

Then you tell LILO to install the boot block using the settings from the mounted partition with:

```
/mnt/sbin/lilo -r /mnt
```

FORGETTING THE ROOT PASSWORD

If you've forgotten your root password, don't panic, as this is easily fixed. Boot into single user mode and run `passwd` to reset the password. If you don't like the idea that resetting the root password is easy, you can tell LILO to ask for a password before allowing you to boot an image.

If you've forgotten your LILO password as well (better cut down on the caffeine!) you'll need to boot into single user mode from your rescue disk, telling it where your root partition is. For example:

```
linux single root=/dev/hda3
```

FILESYSTEM PROBLEMS

If Linux finds major problems with your root filesystem during bootup, it will ask you to run the disc-checking tool `e2fsck` manually as Linux will not try to fix major problems in case it makes the problem worse. If you need to check and repair problems manually, you can do so by typing (using `/dev/hda3` as an example):

```
/sbin/e2fsck /dev/hda3
```

You will be asked various questions about what it should fix, and should answer 'y' to all the questions to repair the problems it finds.

DOCUMENTATION

Most Linux commands are documented in the Linux manual. The system's manual takes a little getting used to, but once you learn its idiosyncrasies it's incredibly easy to find exactly the information you're looking for.

The manual contains at least one page of information for almost every Linux command. The pages of the manual are accessed using the `man` command. For example, to find out about the command `more` type:

```
man more
```

The Linux manual is arranged in sections which apply to different parts of the system. If more than one page exists for a given command, `man [command]` will show only the first man page about the command. If another man page exists, it will



generally be referenced in the 'See Also' section (at the end) of the man page.

Man pages have a title and a number in brackets, for example 'sync (1)'. This means the 'sync' man page is in section 1 of the Linux manual for user commands. Another man page exists in section 2, System Calls, for the system call 'sync'. To view the man page in section 2 you would use:

```
man 2 sync
```

If you aren't sure what command you are looking for, you can use the apropos command. This looks through all the man page descriptions to see if the word given matches any words in the description. For example, the command:

```
apropos filesystem
```

will show a list of all the man pages that are related to filesystems. To print man pages you can just pipe the man command to the printer using:

```
man [subject] | lpr
```

This will print a plain ASCII text version of the page. To get a formatted version of the man page, use:

```
man -t [subject] | lpr
```

The -t switch produces postscript output with bold text and italics in all the right places.

MANUAL SECTIONS

Section	Description
1	User Commands
2	System Calls
3	Library Functions
4	Special Files
5	File Formats
6	Games
7	Miscellaneous
8	Administration and Privileged Commands

INFO

Some commands, particularly those written by the GNU Project, are documented using a system called GNU Info. This system is a hypertext system like the Web for plain text only. To view info pages type:

```
info [command]
```

BOOT AND KERNEL MESSAGES

When your Linux system boots up, a number of interesting messages scroll past, often too fast to read.

KEYS FOR VIEWING INFO PAGES

ARROW KEYS

- [SPACE]
- [ENTER]
- n
- p
- u
- h
- ?

SCROLL AROUND A PAGE

- Scroll forward one page, or move to next page
- Follow the link (menu item) that the cursor is pointing at
- Next page
- Previous page
- Up one level
- Info tutorial
- List of keys

These messages can be very useful, particularly if something has gone wrong during the boot process. You can look at these messages in the /var/log/dmesg file.

Some of the boot messages are not logged in /var/log/dmesg and can only be viewed onscreen. Normally when the system boots and the login prompt is displayed, the screen is cleared. To scroll back and view the kernel messages, you need to stop the login command from clearing the screen. To do this, edit the file /etc/inittab and add the switch '-noclear' to the following line:

```
1:2345:respawn:/sbin/mingetty tty1
```

The next time you boot, the login prompt will appear below the boot messages without clearing the screen first. To scroll back and see the earlier messages, use the Shift-Page Up key combination. If you're booting into runlevel 5, which is the default for Red Hat Linux 8.0 and Mandrake 9.0,

you'll need to first switch to runlevel 3 (console mode) and reboot. To do so, edit the /etc/inittab file and look for the following line:

```
id:5:inittab:
```

Change the 5 to a 3 and save the file. When you reboot you'll be greeted by the console.

PARTITION ORGANISATION

Linux uses a hierarchical filesystem that includes a number of standard directories. These directories contain specific parts of the operating system that may or may not change over time.

Some experienced users give each of these directories their own partition. There's a very good reason for this. It means that, for example, you can upgrade your version of Linux by formatting and reinstalling the root partition without affecting programs you have installed, or the configuration files in your home directory that define how your GUI should look.



If you feel up to dancing the partition shuffle when you next install Linux, try creating a set of partitions and assign them the following mount points:

/

You'll always need your root and swap partitions, of course.

/var

The var directory changes frequently because that's where most log files reside, as well as mail and printer spools, and temporary cache files. Because files under this directory change often you can reduce file fragmentation by keeping this directory on its own partition.

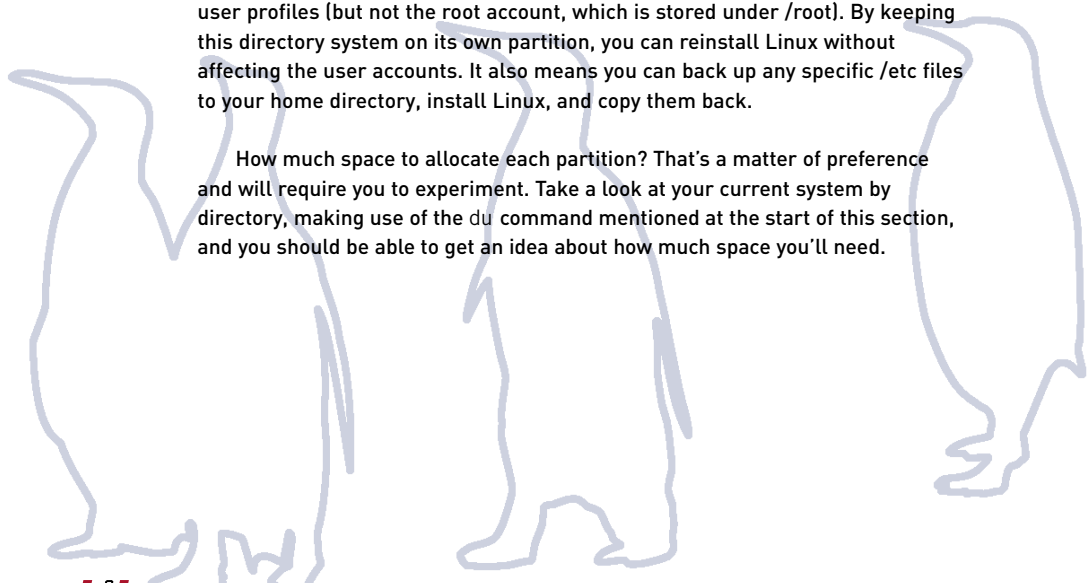
/usr

This is where most programs, libraries, documentation, man pages, and so on are installed. Files change here only when you install or uninstall programs. Sometimes it's also recommended you create a /usr/local partition as this is often the location you install your own applications. This way you can reformat and reinstall the /usr directory during a new installation, but not overwrite the programs you've installed into /usr/local.

/home

This is where all the information relating to users is stored, including your own user profiles (but not the root account, which is stored under /root). By keeping this directory system on its own partition, you can reinstall Linux without affecting the user accounts. It also means you can back up any specific /etc files to your home directory, install Linux, and copy them back.

How much space to allocate each partition? That's a matter of preference and will require you to experiment. Take a look at your current system by directory, making use of the du command mentioned at the start of this section, and you should be able to get an idea about how much space you'll need.



Configuration files **You will, at one point or another, need to configure an option or feature through a configuration file, and when you do it's handy to know where to look.**

To round off this section of the book we're going to show you where some of the most important configuration files are located and what they do.

/etc/profile

The /etc/profile file contains system-wide environment variables. The configurations in this file will affect all users on a system. It includes settings such as the PATH variable.

/etc/motd

The /etc/motd (message of the day) file is used by systems administrators to place a text message that will be broadcast to users when they log on. It can be a useful tool for alerting users to important information, or if you just want to be greeted by your PC when you logon in runlevel 3.

/etc/rc.d/rc.local

The rc.local file contains a bash script that is executed last during the system boot process. If you need to execute a command every time the system boots, you can place that command in here.

/etc/HOSTNAME

The hostname file contains the

complete hostname (including the domain name) of your system. If you need to reconfigure the hostname for any reason, you can edit this file.

/etc/ld.so.conf

The ld.so.conf file lists the directory search order for library files. If you install a program that uses a library that isn't in the default path, you will need to add the path to this file, and then run the ldconf program to initialise the update.

/etc/cron.*

The cron directories hold scripts that are run at specific intervals. There are four different cron directories, which will execute hourly, daily, weekly and monthly. The files inside of the cron directories are shell scripts, so you can easily add your own.

/etc/hosts

The hosts file contains a listing of all known hosts and IP addresses for your machine. If you are running a home network, you will want to add the hostnames and IP addresses of the computers to this file.

/etc/login.defs

This file specifies various system defaults that generally pertain to login definitions. The most important are the password expiration dates and mail configuration.

/etc/inittab

The inittab file controls general options that are available when the system boots or shuts down. The most important option to the majority of users will be the first uncommented line, which specifies which run level the system should boot into. The default runlevel for both Red Hat 8.0 and Mandrake 9.0 is 5 (graphical). If you wish to boot straight to the console rather than booting into X, you can change this to 3 (multiuser).

/etc/resolv.conf

The resolv.conf file is the file where the IP addresses of the primary and secondary nameservers are defined.

/etc/smb.conf

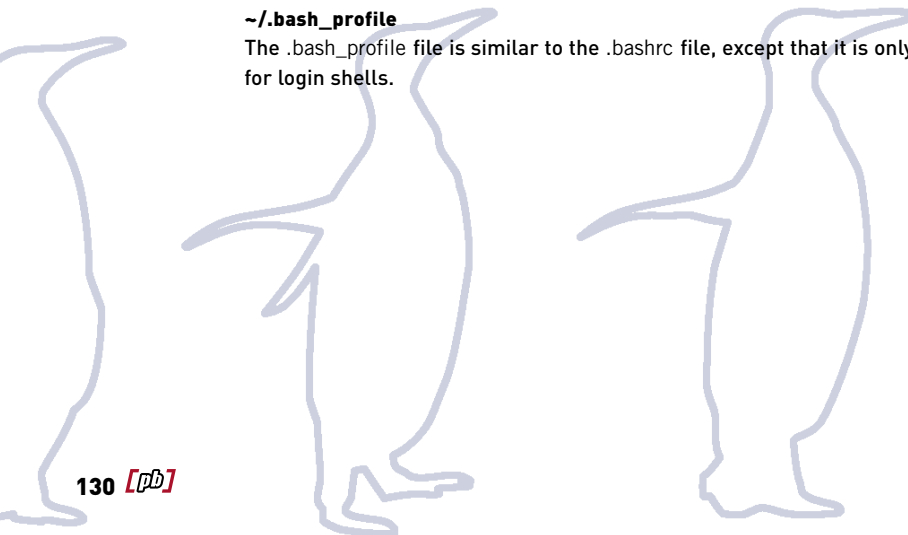
The smb.conf file contains the configuration parameters for Samba. For more information on Samba, see page 138. For Red Hat this file can be found under /etc/samba.

~/.bashrc

The .bashrc file contains information related to the configuration of the bash shell. It is read every time a new bash shell is opened.

~/.bash_profile

The .bash_profile file is similar to the .bashrc file, except that it is only opened for login shells.





Working with Windows

WHAT ABOUT WINDOWS? 132

WINDOWS CONNECTIVITY 138

BASIC SECURITY 145



What about Windows? Inevitably, there will be some Windows programs you simply can't live without. Fortunately, you can run Windows programs under Linux — with a little help.

In the 2nd edition Linux Pocketbook we covered the installation, configuration and use of VMware as a means running Windows under Linux. While VMware has matured into a highly reliable and feature-rich emulator to run almost any operating system under Linux (including Linux itself), it's still not terribly fast. This is not so surprising, given your machine has to think for two.

In this edition of the Linux Pocketbook we're going to show you how to use a program called Win4Lin. Like VMware, Win4Lin is an emulator, but only in so much as it emulates enough of an environment to support Windows and nothing more. The end result is speed. In fact, Win4Lin is so fast at running Windows that Windows installs quicker under Win4Lin than it does on a bare system!

Before we delve into Win4Lin we should point out that there are other options available. Aside from VMware and Win4Lin, you can also use Bochs and Wine, both open source projects, to work with Windows. Bochs has been under development for many years, and is perhaps the first serious emulator to provide a virtual PC under Linux. It is, however, more of an exercise in programming because it's too slow to be used seriously. You can find more information on Bochs at bochs.sourceforge.net. Wine, also still under development, does in fact work quite well depending on what programs you want to run. Unlike VMware and Win4Lin, Wine is Not an Emulator — rather, it's a 'Windows compatibility layer'. Wine will run Windows applications natively on your Linux desktop with minimal resources — it doesn't even require a copy of Windows to run. Wine can be found online at www.winehq.com.

WIN4LIN 4.0 SPECIAL TRIAL VERSION

The helpful people at Netraverse have kindly organised a special offer for all readers of The Linux Pocketbook 2003 edition — a trial version of Win4Lin 4.0. There is currently no trial version of Win4Lin 4.0 available from Netraverse's website, only a link to buy the software. APC Readers get to try it out for 30 days before deciding whether or not to invest in it. So if you've been looking for the perfect program to run Windows under Linux, now you've got it. Have a look at the APC cover CD or www.netraverse.com for more information.

REQUIREMENTS

First and foremost, Win4Lin will install a copy of Windows directly onto your Linux filesystem, so you'll need an original copy of any version of Windows 95 / 98 / ME and enough disk space to support a full Windows installation. Win4Lin does not support Windows 2000 or Windows XP, as these operating systems don't run from a DOS environment. Win4Lin also requires a 2.4.x kernel, which is already installed as part of Red Hat 8.0 and Mandrake 9.0.

On the hardware side, Win4Lin requires as a minimum a Pentium-class processor and 16M of RAM per Windows session. This means that if you have 64M, it will run just fine in addition to all your usual Linux applications.

KERNEL CONFIGURATION

To work its magic Win4Lin needs additional kernel support that is not included in the distribution kernels shipped with Red Hat and Mandrake Linux. There are two methods of adding this support: download an updated kernel specifically for your version of Red Hat or Mandrake from Win4Lin's home page, or patch your current kernel source with the required updates and recompile your kernel. The latter can be a little tricky even if you have experience compiling a kernel, so it's highly recommend that you take the first approach.

If you're using either Red Hat 8.0 or Mandrake 9.0, you'll find patched

kernels on the APC Cover CDs. Even though both Mandrake 9.0 and Red Hat 8.0 are both brand new distros, the Netraverse coders work hard to keep Win4Lin up to date with the latest releases.

INSTALLING WIN4LIN 4.0

The first thing that you should do before installing Win4Lin is read through the installation notes. Mount the cdrom and browse the root directory for a file called INSTALL.html. Double click on it to open it in a browser window.

After reading the installation notes, browse the APC Cover CD and download the Win4Lin archive, and a replacement kernel for your architecture to your /tmp directory. We have provided kernels for i386, i486 and i586 (Pentium) systems for both single CPU and SMP machines.

Next, open up a terminal, su as root, cd to /tmp and run:

```
rpm -ivh [your chosen kernel RPM]
```

Remember that if you press Tab while typing these lines the filenames will be completed for you.

Once installed you need to boot the new kernel before installing Win4Lin, and to do this you need to add another entry to LILO. Change to /etc and open lilo.conf in your favourite editor. Add the following:

```
image=/boot/vmlinuz-2.4.18-17.win4lin
```

```
label=Win4Lin
root=/dev/[your root partition]
read=only
```

If you're not sure what your root partition is, just look at the entry that has the label 'linux'. The new entry you are creating should look identical to the default 'linux' entry, except for changing the label and the name of the kernel being loaded.

Save the file, run `lilo` to install the new kernel and then reboot. When LILO appears, select your new 'Win4Lin' kernel to boot.

PREPARING TO INSTALL WINDOWS

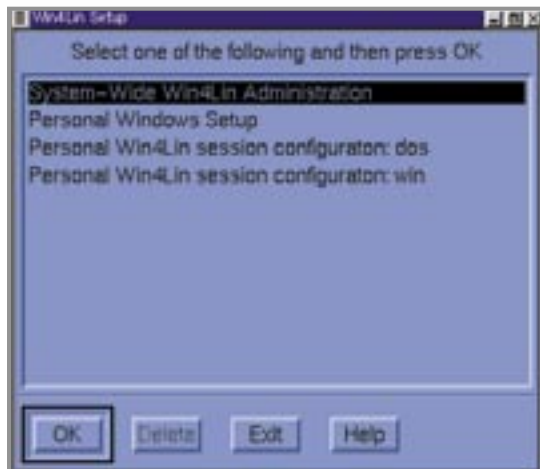
Unmount the APC Cover CD and replace it with your Windows CD. Type the following while still logged in as root: `winsetup &`

The Win4Lin Setup program should appear as shown to the left.

Select the 'System-Wide Win4Lin Administration' option, then 'Load Windows CD'. On the next screen select to install from the CD-ROM. After clicking the 'Install' button, Win4Lin will copy the Windows CAB files to your hard drive ready for an installation.

When that's done exit from the setup program and log out of your root account (type `exit`). Now run `winsetup &` again, this time as your normal user account, and select 'Personal Windows Setup'. Click the 'Install' button and Win4Lin will create a DOS environment from the copied Windows files, boot it and then start the standard Windows installation.

Configuring
Win4Lin.



Win4Lin will perform a scripted Windows install to fit with its environment. The only time you have to interact with the installation procedure is when you type in your Product Key from the 'Certificate of Authenticity' for Windows. Everything else is set up automatically for you.

When Windows has finished installing, select Start → Shut Down in Windows to finish the Win4Lin installation.

USING WIN4LIN

To start a Windows session using Win4Lin, simply run `win &`

Depending on the speed of your



machine, you should find Windows running under Win4Lin to be almost as fast, if not as fast, as running Windows alone. Applications such as Word, Excel and Netscape will all work just fine. The only Windows applications that won't work quite as well are programs designed to perform system-level tasks, such as Disk Defragmenter or networking tools. In order to provide support for certain Windows functions, such as networking, Win4Lin replaces some core files with its own that interact with Linux (you can see these in the installed Windows directory). This means that some networking tools won't work properly — but then you shouldn't need any Windows networking tools as Linux comes with plenty of its own. There's also no need to configure networking in Windows through the Control Panel, as Win4Lin will be directly using your Linux system's settings.

Internet Explorer requires some minor configuration before it will run. Right-click on the IE logo, then select Properties → Connections → Setup. When the Internet Connection Wizard appears click Cancel, select the checkbox 'Do not show the Internet Connection wizard in the future', and then click Yes. This step tells IE not to attempt connecting to the Web through a modem, but instead directly through the network.

Your CD-ROM drive will be automatically mapped to the N: drive under Windows, so you can just put in a CD-ROM and install any applications

you need. Linux and Win4Lin share the CD-ROM perfectly — you can mount or unmount the CD-ROM in either OS and the change will be reflected in both.

LIMITATIONS

Win4Lin will run most applications quite nicely, however some applications will not run because they use features that are not yet supported in Win4Lin. In some cases, these features are supported in VMware, but in other cases you may need to dual boot back to Windows running natively.

DirectX and soundcards are currently completely unsupported, because the focus for this release was desktop productivity applications. If you have any applications that rely on these features, you won't be able to use Win4Lin.

Applications that use Microsoft Networking such as 'Network Neighborhood' will not work. This doesn't include applications that use Winsock networking (which is the majority of networking applications, such as browsers and FTP clients). However, there is a simple workaround you can use to access network machines under Win4Lin.

MAPPING DRIVES WITH WIN4LIN

By default, Win4Lin will make the following drives available under Windows:

A: First floppy drive
B: Second floppy drive
C: Your personal Windows Installation under the 'win' directory



- in your home directory
- J: Set to '/var/win4lin/dosroot', which is where the files for the initial DOS installation are copied to from the Windows CD
- N: CD-ROM drive

To add additional drives to access other locations in your Linux filesystem, you need to run winsetup again. Select 'Personal Win4Lin session configuration: win', click OK, then select 'Drives & Filesystem'.

Click the 'Add...' button to specify the parameters of the drive you want to add. You can specify a drive letter, drive type and location that the data is stored. The new drive mappings will take effect when you restart Windows.

Win4Lin honours the usual Linux permission structure, so sharing, for example, the root directory will give you the same access to the filesystem hierarchy as you would logged in on your user account.

ACCESSING REMOTE WINDOWS SHARES

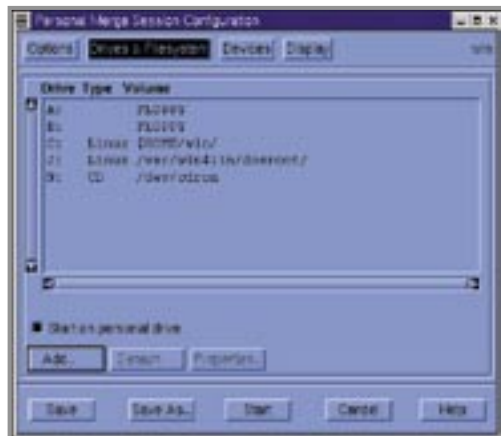
Because Win4Lin does not yet support Windows Networking, accessing remote Windows shares is not as convenient under Win4Lin as it is using native Windows networking, but thanks to Linux's support for mounting SMB shares, it is possible. Samba networking is covered in more depth in the next section, but for now you can mount a remote SMB (Windows Networking) share under Linux, by suing as root and typing:

```
smbmount //[hostname]/[sharename] /mnt/[directory]
```

Hostname is the name of the remote Windows host you are trying to connect to, and sharename is the name of the share. The mount directory should already exist and be empty (if it's not empty, its contents will temporarily disappear until you unmount the remote share).

When smbmount prompts for a password, type in the password for the share, or just press enter if it's a public share. With the SMB share now connected to Linux you can simply map the directory into Win4Lin to be available under Windows by following the instructions on the previous page.

Drive mappings under Win4Lin.



If you don't want to have to su as root every time you want to mount a Windows share, you can allow any user on your system to mount SMB shares by typing:

```
chmod u+s /usr/bin/smbmnt /usr/bin/smbumount
```

UNINSTALLING WIN4LIN

If you decide that Win4Lin is not for you, then you can remove it by running each of the following sets of commands, which undo the steps you performed at installation. Of course, you'll need to be careful when running these commands to ensure you don't delete anything you may require. To remove Win4Lin simply type the following:

```
rpm -e Win4Lin
```

To remove your personal Windows installation:

```
cd ~
rm -rf win
```

To remove the system Windows setup:

```
cd /var
rm -rf win4lin
```

Be sure to backup any important information you've created whilst using Win4Lin first!



Running Windows under Linux with Win4Lin.



Windows connectivity

There will be times when it's handy to be able to share files and services between Linux and Windows machines. Enter the joy that is Samba.

Samba is a software package included with Red Hat and Mandrake Linux which allows you to provide file sharing and print services to SMB clients. SMB is the Server Message Block protocol, also known as CIFS or Microsoft Networking. Windows for Workgroups, Windows 95, 98, NT and 2000 all use SMB as their native file and print sharing protocol, so Samba is the ideal product to turn your Linux machine into a file server for Windows clients on a network.

Samba can also act as a WINS (Windows Internet Naming Service) server, a Windows Login Server (for Windows 95 and 98), and a Windows Domain Controller, although Domain Controller support is still experimental.

INSTALLING SAMBA

Depending on the options you selected when you installed Linux, you may already have Samba installed. To check from a console type:

```
rpm -qa | grep samba
```

If 'samba', 'samba-common', and 'samba-client' are all displayed, then you can move on and learn how to configure Samba. If you only plan to use shares exported by other Windows machines on your network, and not provide accessible shares from your Linux box, the required Samba components (samba-common and samba-client) should already be installed, so you can skip to 'Mounting Windows resources' to learn how to access these resources. You only need the 'samba' package to export shares for other machines to use over the network which, we'll show you how to configure here.

If one or more of these packages is missing you'll first need to install them from the Red Hat or Mandrake Linux CDs (look for the RPMS directories). For help with installing applications, see page 113.

CONFIGURING SAMBA

Nearly all Samba configuration is performed through the file `smb.conf`. It's a large file with many options, but there are only a couple of settings that need to be

changed to configure Samba for your system.

There is a nice, easy graphical front end for configuring Samba called Swat (Samba Web Administration Tool), but we're going to show you how to use this after we've shown you how to configure Samba the traditional way.

There's a good reason for this; Swat won't always be available to you, so it's very handy to know how to quickly and easily configure Samba from the command line. Read ahead if you want to, but don't say we didn't show you how to do it the old-fashioned way!

Login as root and change over to /etc/samba under Red Hat or /etc under Mandrake and open `smb.conf` in your favourite editor.

1 STEP ONE

WORKGROUP

The first item to configure is the workgroup to which your machine belongs; this needs to be the same as the workgroup setting on your Windows machines. The workgroup setting is near the top of the `smb.conf` file, and will look something like:

```
workgroup = Pocketbooks
```

If, later, you can't find your Linux machine when browsing the network from a Windows machine, check again that you've set the workgroup correctly.

2 STEP TWO

ACCESS

If your computer or network is connected to the Internet, you will almost certainly want to configure the 'hosts allow' option, so that only hosts on your local network are allowed to connect to your Samba server. For example, if your local network is 192.168.0.0, you might edit the following line as such:

```
hosts allow = 192.168.0. 127.
```

Note how a class C network is specified by '192.168.0.' (with a trailing period). The '127.' entry is added so that you can test the Samba server from your machine. Yes, that is a space between '0.' and '127.'

3 STEP THREE

SECURITY

If you want Samba to behave like a standard Windows machine you'll need to set the security access to 'share'. This will enable other machines to access the resources on the Samba machine without requiring any login or password. Look for the 'security' entry a little further down the file and set it to:

```
security = share
```

4 STEP FOUR

SHARING

Now you need to create a new share



to make available to the Windows world. Scroll down to the bottom of the file until you find an area headed 'Share Definitions'. You will see a variety of preconfigured, commented-out share definitions to serve different purposes. Scroll down to the bottom and you'll find one for 'public'. Simply uncomment this (remove the preceding ';' from each line) and edit the settings accordingly — set the correct path, and change the share name from [public] if you wish — to quickly and easily set up a basic public share that everyone can access. A public share might look like this:

```
[public]
  path = /usr/downloads
  public = yes
  only guest = yes
  writable = yes
  printable = no
```

You will see that each of the example shares are clearly documented. Spend some time reading the examples to learn about the different types of shares you can create. For public shares you need make sure 'public' is set to 'yes' and, of course, if you want anonymous users to have write access, to set 'writeable' to 'yes'.

5 STEP FIVE

TESTING

After saving the file, you should run the `testparm` program which will read through the configuration file and check that you didn't make any syntactic errors. The next step is to start Samba (or restart it, so that it reads its configuration file again) by running:

```
/sbin/service smb restart
```

To ensure that Samba starts the next time your system boots, you'll also need to run `ntsysv` and make sure 'smb' is selected. The `ntsysv` program is covered in more depth in the 'Basic security' section of this Pocketbook coming up next.

Before going to a Windows machine to test Samba, it's a good idea to first check it locally with `smbclient`. To list the shares available, you can use the `-L` switch and specify the server to connect to:

```
smbclient -L localhost
```

If you are prompted for a password, just press Enter without typing anything.

You should at least see an `IPC$` and your new share listed.

The next step is to try and connect to the share you've made from a Windows machine across the network. First make sure your network is properly set up and working (you can ping other machines). If you can't find the Linux machine under the work-group you've set up under Network Neighborhood, click on Start → Find → Computer and type in the IP address of your Linux box.

MORE ON SECURITY MODES

There are two main security modes that an SMB server can run in: 'share level security' and 'user level security'. With user level security (`security = user`), a username and password are required to gain access.

Generally, if you login to your Windows computers with usernames that correspond to those on your Linux server, then you would choose user level security, otherwise you can choose share level security.

ENCRYPTED PASSWORDS

If you've decided to use `security=user` there's one more step you may need to perform to make your shares available. All recent versions of Windows (including Windows 98 and Windows NT 4.0 Service Pack 3 or later) insist on using encrypted passwords when communicating with SMB servers. Unfortunately, the encrypted password scheme that is built into the SMB protocol is incompatible with the way that Linux

usually encrypts passwords.

There are two ways to get around the encrypted password problem. One way is to adjust the Windows client configuration so that Windows no longer requires encrypted passwords. The second way is to configure encrypted password support into Samba.

To tell Windows to ignore encrypted passwords, you need to change a particular registry setting. To assist you in changing it, the Samba documentation includes a REG file (one for each version of Windows); simply copy it to the Windows machine, then double-click on it to integrate the change automatically in the Windows registry. The registry files are located in `/usr/share/doc/samba-[version]/docs`, and are named, for example, 'Win95_PlainPassword.reg'.

Configuring Samba to support encrypted passwords is a little more involved. You first need to create a file called `smbpasswd` in the directory

Editing
smb.conf.





where `smb.conf` resides by typing (as root):

```
mksmbpasswd.sh < /etc/passwd > smbpasswd
chmod 600 smbpasswd
```

This `smbpasswd` file will initially include all the logins currently in your system (as defined in `/etc/passwd`), and will be used to store the SMB-encrypted passwords.

The next step is to set the passwords in the file by running `smbpasswd`. To set the SMB password for the user 'pocketbooks' you would type:

```
smbpasswd pocketbooks
```

After defining the SMB-encrypted passwords, you need to tell Samba that you want to use encrypted passwords by uncommenting the following two lines in `smb.conf`:

```
encrypt passwords = yes
smb passwd file = /etc/samba/smbpasswd
```

If you're using Mandrake, use `/etc/smbpasswd`.

If you restart Samba, encrypted password support will be enabled. Note that the SMB-encrypted password and the regular Linux password are stored in two separate files and are independent from each other. If, when you add new users to your Linux system, you want them to have access through Samba, you shouldn't run `mksmbpasswd.sh` as you initially did, as that will delete any encrypted passwords you have since added. Instead, you should run `smbadduser`:

```
smbadduser apc:apc
```

You need to specify the username twice, as it is possible (although perhaps confusing) to have different Linux and SMB usernames.

MOUNTING WINDOWS RESOURCES

We've covered how to share resources on a Linux system and how to make them available to Windows machines on a network, but what about the other way around? Samba is also used to allow Linux machines to mount shared Windows resources onto its own filesystem tree across a network.

The command to do this is `smbmount` (although you can also use the `mount` command). Mounting a network share is very much the same as mounting a local drive. First `su` as root and then create a directory where you intend to mount the share. For example:

```
mkdir /mnt/share
```

Assuming you want to access a shared drive called 'downloads' on a Windows machine at '192.168.0.10', you can then mount the network share as follows:

```
smbmount //192.168.0.10/downloads /mnt/share
```

You will be prompted for a password. If it's a public share, just press Enter and the share will be mounted. If it's a password-protected share, typing in the password will give you access. If you know it's not a protected share you can bypass the password prompt by appending '-o guest' to the command. When you've finished accessing the remote share, you can unmount it by typing `smbumount /mnt/share`.

If you don't want to have to `su` as root every time you want to mount a share, you can allow any user on your system to mount shares by typing:

```
chmod u+s /usr/bin/smbmnt /usr/bin/smbumount
```

If you want to permanently attach an SMB share every time your Linux machine boots up you can simply add the SMB mount to your `/etc/fstab` file:

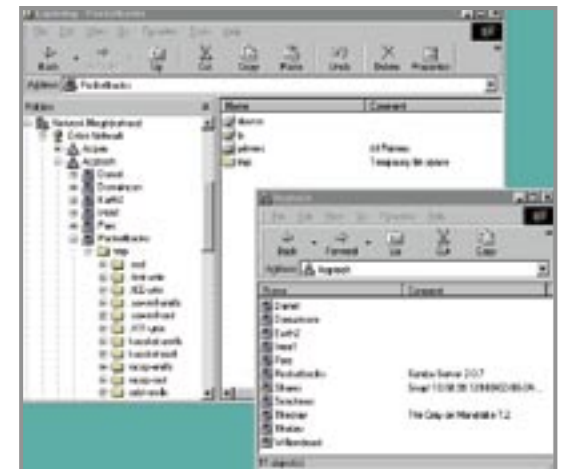
```
//[host address]/[share name] /mnt/[mount point] smbfs guest 0 0
```

This will cause the directory to be automatically mounted at boot time using the 'guest' option to avoid being prompted for a password. If you want to mount the share only on demand, add `,noauto` after 'guest'. Then, when you want to mount the share, you can simply type `mount [mount point]` from the `/mnt` directory to automatically mount the network share.

PRINTING

The default Samba configuration will automatically set up any printers that have been defined in your `/etc/printcap` file. This means that once you've set up printing from within Linux, your printer(s) will automatically be available to any of your SMB clients.

Browsing a Linux box with Windows networking.





SWAT

Swat is a very easy and portable method of configuring Samba. Using a browser interface, Swat allows you to manage all the options of your `smb.conf` file locally or remotely over a network or the Web.

Swat is included with Samba, but isn't enabled by default. In both Red Hat 8.0 and Mandrake 9.0 you can enable it by editing the `/etc/xinetd.d/swat` file and changing the line which reads 'disable = yes' to 'disable = no'

For security reasons, this default configuration will only allow Swat to be accessed from the localhost, but by commenting out or editing the 'only_from' line, you will be able to access Swat from any location from a Web browser.

After saving the file, run `/sbin/service xinetd restart` to implement your changes.

Once enabled, you can access Swat by simply pointing your Web browser to `http://127.0.0.1:901`. You will be prompted for a username and password; simply login as root. Port 901 is the port Swat uses, as defined in the `/etc/services` file. To use Swat remotely just connect using the IP of your Linux machine. Note: although Swat is Web based you don't need Apache running to use it.

Using Swat is straightforward. If you need help with Samba you can easily access comprehensive documentation from within Swat. Each of Samba's configuration sections, including managing shares, can be customised using the buttons along the top of the page. The Status page shows (as you'd expect) the current status of Samba and also allows you to start, stop and restart the Samba servers `smbd` and `nmbd` — all from the comfort of your

browser. You can also see who is connected and who is using Samba shares and even, if you feel the inclination, kill access by connection. The View page will show you the `smb.conf` file which it has generated for you, and the Password page lets you add and remove users as well as change passwords.

FURTHER INFORMATION

Samba is well documented. In addition to visiting the Samba home page at samba.anu.edu.au/samba, trusty `man smb.conf` will provide plenty of helpful information.

Configuring Samba using Swat.



Basic security

No matter what operating system you use, your machine is always vulnerable to crackers when connected to the Internet. The following guide will help you secure your system when online.

The purpose of this section is to provide some essential information that will help you protect your Linux system from intruders when connected to the Internet.

Some people think that if their system is only intermittently connected to the Internet by modem that they are safe, but this is certainly not true. Crackers running automated host scanning tools won't know that you aren't permanently connected until you log off.

There is also a prevalent belief that security isn't important for some people because their machine doesn't have anything important on it. But if that machine is on a LAN then compromising other hosts on that LAN will be far easier once even one machine has been compromised. In addition, a compromised machine is often used as a stepping stone to attack other hosts so the true origin of an attack is disguised.

That said, it's possible that your machine may never be the target of an attack, but the importance of that 'may' is equal to the importance of the information you keep on your machine or local network. If you plan

to provide a permanent connection to your Linux box, such as through cable, then you should think seriously about ensuring it is adequately secured against unwanted guests.

There are three main principles of network security that are important to follow:

- ✦ Disabling or removing unneeded network services
- ✦ Keeping up to date with current security updates
- ✦ Restricting access to services and granting access only to hosts that require it

DISABLING UNUSED NETWORK SERVICES

The idea behind disabling or removing unneeded network services is that you cannot be affected by bugs in software that you are not running. For example, if a bug is discovered in a network service that would permit users without a local account to gain root access to your system, that would only be a problem if you are running that service.

Simple as this may be, you'd be surprised by how many machines



connected to the Internet are running network services (FTP, Telnet, HTTP, and so on) that the owners of the machines themselves never make use of.

Logically then, the fewer services you run, the less chance you will be caught unawares by any as-yet-undiscovered bugs. By removing any unused services you immediately reduce the risk of your system being compromised.

If you performed an install that didn't include many networking services, then it may be that some or all of the services won't be present on your system. If this is the case, don't worry — this is a good thing from a security point of view. The purpose of this section is to explain how to disable services. If they aren't installed, then you don't have to worry about them.

To disable or remove services, you first need to know how the various services are started. Network services fall into two main categories: those started standalone, and those started through inetd (recently evolved to xinetd in Red Hat 8.0 and Mandrake 9.0).

STANDALONE NETWORK SERVICES

Services started as 'standalone' are started when your system boots, and then continually wait for any new incoming connections. Each of these services has a script in the `/etc/rc.d/init.d` directory which is used to control them. You can run the scripts to start, stop or check the status of a service. For example, from the `/etc/rc.d/init.d` directory you can type:

```
./smb status
```

to see if Samba is currently running. For a list of options for the script, you can just type `./smb`. You can also run `/sbin/service smb status` to perform the same task.

These scripts will allow you to manage services on the fly, but what about determining which services are started automatically at bootup? There are many

ways to do this but the easiest is by using a handy tool included in both Red Hat and Mandrake Linux — `ntsysv`.

NTSYSV

After running `/usr/sbin/ntsysv` from the command line, you will be presented with a list of services set to run during bootup. You can scroll up and down with the arrow keys and select or deselect services with the space bar. To find out what a service does, press F1 to read the description. When you've

ntsysv can be used to select startup services.



A NOTE ON SENDMAIL

If your email address resides with your ISP (which it will, unless you're running your own mail server on the Web), you might be tempted to disable sendmail (or postfix if you have it installed). This is not a good idea, as some system management programs will email the root account to inform you (as system administrator) of anything important that happens with the system.

By editing `/etc/sysconfig/sendmail` you can prevent sendmail accepting any incoming communication while still allowing it to deliver any mail to local accounts generated by local processes. The line to look for in that file is `DAEMON=yes`, which you should change to `DAEMON=no`.

After modifying `/etc/sysconfig/sendmail` you will need to restart sendmail by running `/sbin/service sendmail restart`. Alternatively, you can pretend you're using Windows and force a reboot to initialise the changes.

finished selecting and deselecting services, Tab across to 'OK' and your changes will be saved.

In the past, core services such as FTP and Telnet were only managed through inetd, but both Red Hat and Mandrake have simplified this by making these services available in programs like `ntsysv`. This is a welcome change as it makes it easier to manage the starting and stopping of all essential services. But be aware that these services are still only executed on demand and are managed through `inetd` and `xinetd`.

Which services should you disable? Depending on the type of install you performed and the packages you selected, you might not need to disable any services. Similarly, what you have planned for your Linux box will determine which services you need. If you just plan to surf the Web and use Linux as a desktop OS, you can safely disable:

httpd	Apache Web server
telnet	Telnet server
wu-ftp	FTP server
inetd/xinetd	Internet super-server (see notes below)

If you're not sure what a service does, even after reading Help, leave it on its default setting.

When you exit `ntsysv`, it will not stop or start any services by itself. If you want a service to be immediately disabled or enabled (rather than waiting until your next reboot), you need to run:

```
/sbin/service [servicename] [start/stop]
```

NETWORK SERVICES STARTED FROM INETD/XINETD

`inetd` is the Internet super-server daemon and is responsible for starting specified Internet services on



demand. This means, for example, the FTP server is only started when someone actually connects to download or upload files.

If you intend to just use your Linux box as a desktop machine to surf the Web, and have no need for services such as the Telnet and FTP servers, you should disable them to stop others using them as possible gateways into your machine. Even though both Telnet and FTP require authenticated logins for people to get in, rest assured that the art of cracking is very much alive. Remember the first golden rule above: if you don't need it, disable it — it's one less handle for crackers to use.

In the past, all services run through `inetd` were configured through the `/etc/inetd.conf` file. To disable any service, you needed to edit this file, comment out the line corresponding to the service you wanted to disable, and then reload the daemon.

Recently `inetd` has evolved into `xinetd`. For those who are familiar with `inetd`, the most obvious difference is that there is now one configuration file per service (stored in the `/etc/xinetd.d` directory) rather than one configuration file for all services.

To disable a service that is run through `xinetd`, edit the corresponding file in `/etc/xinetd.d` and add a line reading `disable = yes` within the braces.

Disabling Telnet and FTP is a good idea if you will be connecting your machine to the Internet and don't plan to use these services. So, for example, to disable Telnet you edit the appropriate entry in `/etc/xinetd.d/telnet` as such:

```
service telnet
{
    disable = yes
    flags      = REUSE
    socket_type = stream
    wait       = no
    user       = root
    server     = /usr/sbin/in.telnetd
    log_on_failure += USERID
}
```

You can then type `/sbin/service xinetd restart` so that `xinetd` will pick up the changes you've made. Alternatively, using `ntsysv` will do this for you for most services.

For a machine that is only used as a workstation, there are no absolutely essential services that are run from the Internet super-server (which is not true of some standalone services), so it's possible to completely disable `xinetd` or `inetd` (as it is itself a standalone service) from within `ntsysv`.

COMMON NETWORK SERVICES

SERVICE NAME	DESCRIPTION
apache	The Apache Web server
fetchmail	Only required if you want to manage email locally
finger	Displays information about users on request
gated	The routing daemon
identd	Reports the identity of users connecting to your system
ipop2/ipop3/imap/imap	Allows users to remotely access their mail on your system
inetd/xinetd	Internet super-server, manages Internet services
ipchains	Sets up firewalling rules
linuxconf-web	Starts the linuxconf-web service
mars	Netware compatible file and print server
mcserv	Service that allows remote file management using Midnight Commander
named	Domain name server
netfs	Automatically mounts remote NFS, SMB and NCP services as defined in <code>/etc/fstab</code>
nfs, nfslock, portmap	Used to share files with other Linux machines using the Network File System (NFS) protocol
pppoe	PPP-over-Ethernet driver used for some ADSL connections
rexec, rlogin, rsh	Remote shell and login
sendmail/postfix	The popular mail servers
smb	The Samba SMB server
snmpd	Simple network management protocol
swat	Remote Samba administration
talk/ntalk	Service to allow users to chat with each other on local/remote systems
telnet	Remote login service
webmin	Remote system administration
wu-ftpd/proftpd	FTP services

The table above shows some common network-related services that are run either standalone or from the super-server. Unless you require them, disable them. If you're not sure whether you need a particular service or not, leave it enabled.

STAYING UP TO DATE

Unless your machine is not connected to a network at all, it's a good idea to stay current with any patches or updates that are released by Red Hat or MandrakeSoft. If security-related problems are found

in packages they have released, then an update will be released as soon as the problem has been fixed (usually within a day or two).

Red Hat's security advisories and updates Web page is located at www.redhat.com/apps/support/updates.html and MandrakeSoft's is located at www.linux-mandrake.com/en/security.

To receive notification by email when new packages are released see www.redhat.com/mailling-lists for information about the redhat-watch-list, or www.linux-mandrake.com/en/flists.php3 for information about MandrakeSoft's security-announce list.

Updates for both distributions are locally available by FTP from <ftp.aarnet.edu.au> in the `/pub/redhat/updates` or the `/pub/mandrake/updates` directory.

LEARNING MORE

There are many topics that haven't been covered in this section on security. For example, if your computer needs to act as a server, you may want to restrict which hosts can connect to what services (the third principle of network security, mentioned at the start of this section). Security will be covered in much greater depth in *The Advanced Linux Pocketbook*.





Playing with Linux

LINUX MULTIMEDIA 152

MULTIMEDIA PLAYBACK 156

GAMES 160



Linux multimedia

Like many facets of Linux, multimedia support is constantly under development and there are currently plenty of options for making the most of multimedia with Linux.

Linux has never been the platform of choice for multimedia applications, but things are starting to change. There have been two main issues preventing Linux multimedia from developing sooner. First, much of the basic infrastructure for supporting high-end multimedia hasn't so far existed and, second, the lack of drivers for cutting-edge hardware. Massive inroads have been made into both these areas, but there's still some way to go. However, the speed at which things are changing is astonishing, and it's only a matter of time before Linux catches up with Windows in the one area where Windows is a recognised performer.

XFREE86 AND MULTIMEDIA

As discussed in Chapter 3, XFree86 is the Linux network graphics protocol. It not only allows Linux applications to show themselves onscreen, but also on other machines across the network via a special language called the X protocol. While X is a nifty idea for business applications, it does slow things down, as information must be converted into the X protocol before it is displayed onscreen. This has an impact on multimedia applications, as speed is often paramount for multimedia applications.

The newest release of XFree86, version 4.0, supports a new way of talking to the screen called Direct Rendering Infrastructure. Its name is an accurate description of its function as it allows applications to talk to the video card, bypassing the X protocol. XFree86 4.0 also has many other additional features which make it better suited for multimedia than the 3.x series. Both Red Hat and Mandrake come with the 3.3.6 and 4.0.1 versions of XFree86, and install the appropriate version based on your hardware (3.3.6 comes with support for a wide variety of older hardware; 4.0.1 supports newer cards, especially those with 3D acceleration).

3D HARDWARE SUPPORT

Soon after the release of XFree86 4.0, many hardware manufacturers that were late in embracing Linux, released Linux drivers to work with the new

architecture. As a result, Linux support for current 3D chipsets is near ubiquitous. The performance of some of the more recent XFree86 drivers has been quite impressive, especially nVidia's GeForce and GeForce2 drivers (a direct port of their Windows drivers) which generally perform at the same level as their Windows counterparts.

MATROX

Matrox has drivers for G200, G400 and G450 cards (and upcoming drivers for the soon-to-be-released G800). Matrox drivers are provided by the Utah GLX project, an open source effort to support major 3D cards in XFree86 3.x. The drivers generally work quite well, and with some performance tuning can eclipse their Windows counterparts. The group of open source developers that work on the Utah GLX project includes John Carmack, the man behind id Software's Quake games. Matrox cards, while not performing under Linux quite as well as their Windows counterparts, are still solid performers, and the dual-headed feature on some cards makes them an attractive proposition for Linux gamers.

ATI

Linux drivers exist for the Rage Pro and Rage 128 cards. Both Rage cards also use the Utah GLX drivers and XFree86 3.x in their 3D support. The Rage 128 and 128 Pro drivers, and the upcoming Radeon drivers (due for release by the time you read this) are

based around XFree86 4.0. They can be downloaded directly from ATI.

NVIDIA

nVidia was a late entrant into the Linux realm, but was certainly worth waiting for. It has released unified drivers that support its full range of cards, including the TNT1/2 and GeForce 1/2/MX/Ultra. nVidia cards uphold their popular reputation with an excellent set of Linux drivers that are mainly on a par with Windows. nVidia cards are fast becoming the card of choice for Linux users everywhere.

The performance of some of the more recent XFree86 drivers has been quite impressive.

3DFX

One of the longest supporters of Linux, with drivers for the Rush, Banshee and Voodoo1, 2 and 3 cards, as well as beta Voodoo4 and 5 drivers. 3dfx's current drivers, while certainly more than capable of handling the current generation of Linux games, lack the speed of their Windows counterparts, lagging some 25% behind in performance on a Voodoo3. On the upside, 3dfx's drivers are extremely easy to use and the linux.3dfx.com site provides an excellent set of drivers. Old Linux games, such as Myth II: Soulblighter, require 3dfx cards due to their reliance on the 3dfx exclusive Glide rendering system. Voodoo1, Voodoo2 and Voodoo Rush use XFree86 3.3.x and the Glide-based fullscreen drivers that 3dfx maintains.

Voodoo3, 4 and 5 use XFree86 4.0 drivers (included with XFree86 4.0.1), but, once again, there are performance issues. The beta generation of Voodoo4 and 5 drivers currently only supports one of the two VSA-100 chips on board the Voodoo4 and 5, dropping performance from these particular cards by another 40%. Ouch.

However, it does seem that 3dfx is listening to its customers and is producing a brand new, completely rewritten set of drivers for its upcoming cards, which the company claims will be on a par with Windows.

OTHER CARDS

Other 3D accelerator cards known to work under Linux include the Intel i810, Sis6236, and the S3 Virge and Savage3D. All of these cards use the Utah GLX project and XFree86 3.x for support.

APPLICATION PROGRAM INTERFACES

MESA3D AND OPENGL

First person shooters aren't the only games that use 3D. Support for 3D acceleration has now become a mainstream part of many games, including 2D games which use the technology to accelerate features such as transparency. OpenGL is the most popular way of providing these functions in 3D games on Windows, PlayStation and PlayStation2, Nintendo 64, Dreamcast, Macintosh, BeOS, Unix and Linux. OpenGL is a semi-open standard, created for high-end 3D workstations by the graphics masters at Silicon Graphics (now SGI). It's also massively popular; it's the technology behind nearly all of the top-selling 3D games of all time on all platforms. Many implementations exist, but the most popular OpenGL-compatible implementation on Linux is called Mesa3D. As a result, Mesa has become the basis for most 3D applications on Linux.

DIRECTX

DirectX is a set of closed multimedia application program interface (APIs) and components created by Microsoft, and limited to the Windows and Xbox platforms. Different APIs exist for different aspects of multimedia — DirectDraw, DirectSound, Direct3D, and so on. Most Windows games use these interfaces (with the exception of Direct3D, which is often supplanted by OpenGL).

Until about a year ago, implementing DirectX-based applications on Linux was a fairly difficult task. However, Loki Entertainment Software, the Linux gaming house, developed an open solution called Simple Directmedia Layer (or SDL) which acts as a multiplatform implementation of a similar series of APIs. SDL has since become the foundation of many multimedia applications on the Linux platform, and has been used successfully by Loki to convert many DirectX-based commercial games to Linux. We'll come back to Loki and the games that have been ported later in this chapter.



introduction | installing | exploring | using | learning | working | **playing** | resources



WINE

Wine, mentioned briefly in the previous chapter, also forms a small part of multimedia (and gaming) on Linux. For those times when there isn't any Linux support available due to closed file formats or Windows-specific games, Wine comes in very handy. The Wine project has made some impressive inroads in recent releases, with OpenGL and partial Direct X support.

There is a wide variety of games which work quite well under Wine. Check out the Wine Success Stories page at www.linuxgames.com/wine for some suggestions and games that are known to work. The list, however, isn't as up to date as it once was as more manufacturers are releasing games specifically for Linux.

The situation with closed file formats is another issue altogether. To play non-standard Windows files on Linux, all you need to run is a small, simple program called a 'codec'. Convincing certain companies based in Redmond, Washington in the US to produce Linux applications isn't the easiest of tasks. Because Wine is open source, some programmers have produced customised codecs designed only to run a particular Windows application. This approach has worked quite well, resulting in a beta quality method of playing Windows media and other Windows-specific file formats on Linux.

Tribes 2, one of the most anticipated games in 2001, will be released for Linux at the same time as Windows.



Multimedia playback

There are quite a few tools available for Linux to play the popular multimedia formats. Here is a selection of the best.

A major issue with multimedia on Linux has been closed file formats. Despite the wild success of open standards and their use in many large-scale projects (for example, the Internet), the rush for control of the Internet and the streaming media market has resulted in many groups pushing their own closed or slightly modified open standard media formats upon the world. This situation seems to be changing with the finalised MPEG-4 specification making its way into a wide variety of products, but Linux users certainly aren't out of the woods yet.

The past year has shown that Linux developers are determined to develop solutions for closed file formats, with or without the support of the companies that control them. Many companies, like RealNetworks and Macromedia, have provided solutions for playback of file formats

RealPlayer for Linux.



such as RealMedia and Flash. Other companies less supportive of Linux, such as Microsoft, have had their file formats reverse engineered by open source developers with the intention of creating compatible Linux versions. Although several of the applications covered here are still under development, they are all usable and provide an excellent demonstration of what's to come in the future. Of course, that doesn't mean you shouldn't email your favourite multimedia company with a polite note asking for a Linux version of its player!

MULTIMEDIA APPLICATIONS

REALPLAYER

The official RealNetworks RealPlayer is available for Linux with support for



XMMS has a variety of spectacular skins available.

the latest RealMedia format files. A direct port of the Windows version, RealPlayer for Linux has all of the features of its counterpart — with the joyful exception of some of the advertising and 'phone home' aspects of the software. Because Linux is a popular server OS, RealNetworks has also ported RealServer to Red Hat and other major Linux distributions.

You'll find the latest version of RealPlayer, 8.0 at www.realplayer.com.

ADOBE ACROBAT READER

Not strictly multimedia, but certainly one of the most popular programs for any operating system, Acrobat Reader allows you to read Adobe Portable Document Format (PDF) files

— a versatile electronic publishing format. Many companies release white papers, manuals, technical documents, brochures and more using PDF files to allow users to view material the way it was intended.

You'll find the latest version of Adobe Acrobat Reader, at www.adobe.com.

XMMS

There are a wide variety of easy-to-use MP3 players for Linux. Perhaps the most popular is XMMS, an open source clone of Winamp. XMMS has the ability to use skins from Winamp and a variety of other players. It boasts an excellent plug-in architecture, some extremely good-looking 2D and 3D visualisation effects, and superb high-

quality playback. (For encoding MP3s, Grip is an excellent all-in-one ripping and encoding application, and comes bundled with Mandrake 9.0.

XMMS is included in both Red Hat and Mandrake Linux distributions.

XMPS

A player for the controversial DivX video format, which uses a small portion of Wine to run a Windows DivX DLL file to play movies. This approach works surprisingly well.

Download the latest version of XMPS from www.xmps.sourceforge.net.

AKTION

A basic MPEG-1 player built around KDE and Qt, Aktion plays MPEG-1, Cinepak, older versions of Intel Indeo, MJPG QuickTime, and more.

Aktion is included in both the Red Hat and Mandrake Linux distributions.

AVIFILE

Though rumours do circulate every now and then, it's doubtful that Microsoft will create a Linux version of Windows Media Player anytime soon. However, Linux enthusiasts can use AviFile, an open source application that uses a small subsection of Wine to act as an interface for Windows multimedia codecs for Linux. While AviFile is in beta at the time of writing, it's quite stable and entirely

capable of playing many closed file formats under Linux, including Microsoft formats.

Download AviFile from divx.euro.ru.

QUICKTIME FOR LINUX

Quicktime itself isn't a specific codec; it's a standardised 'wrapper' for a variety of codecs, including Sorenson, MJPA and more. While it's possible to view and create MJPA-encoded Quicktime movies on Linux using an open source application called Quicktime for Linux, there currently exists no solution for either playing or creating the more popular Sorenson-encoded media. A future version of AviFile may allow this to be possible, but don't hold your breath. Apple isn't



introduction | installing | exploring | using | learning | working | **playing** | resources

entirely averse to Linux, so a polite email asking for a Linux Quicktime player might not go unnoticed.

Download Quicktime for Linux from heroine.linuxave.net/quicktime.html.

OPEN MEDIA SYSTEM DVD PLAYER

Open Media System is a controversial DVD player for Linux. To understand why, we first need to learn a little about DVDs and their distribution format.

The information on DVDs is stored in an encrypted format to try to prevent piracy. However, DVDs can still be easily duplicated while encrypted, so the furore over DeCSS (the code that allows you to crack the DVD encryption) on the Internet is more about rights of access than piracy.

What this means is that Linux, an open source OS, has been unable to play DVDs since their inception, because any DVD player made for Linux under the Open Source model would need to have distributable source code, and thus the algorithms to decrypt DVDs. When DeCSS arrived it allowed programmers to write a DVD player for Linux, hence the Open Media System.

The MPAA has been trying hard to eradicate DeCSS from the face of the planet, claiming it is conducive to piracy. But, as mentioned previously, it is still possible to copy DVDs in their encrypted format. DeCSS and Open Media System simply allow users to play DVDs on their systems,



not copy them, which can be done without the use of DeCSS.

Unfortunately the powers that be got the better of Open Media System in the end, and development was stopped in 2001. You can still download it from www.linuxvideo.org, but it is neither supported nor worked on by the initial designer. Even though development was frantic at one time, the Open Media System DVD player is proof positive that successful Linux projects require diligent leadership.

SONIQUE 2.0

The next generation of this Windows-based multimedia player is coming to Linux. Sonique is perhaps the best-looking multimedia player on any platform, with excellent skins and animation. The Linux version will play a large range of file formats, including MP3 and possibly also Windows Media.

Download the Sonique player at www.sonique.com.

Open Media System DVD playback on Linux. Oh, and Buffy.

XMPS
and Keanu
in action.





Games

What would a computer system be without games? With speed and stability on its side, Linux is slowly becoming the operating system of choice for gaming.

There's a wide variety of games available from both the open and closed source communities for Linux. Although Linux has some way to go before being as ubiquitous a gaming platform as Windows, the variety of titles is increasing at a considerable rate. This is primarily due to the maturity of the Linux multimedia infrastructure with SDL and XFree86 4.0, and the tendency of many Linux enthusiasts to hassle game developers into compliance!

Linux games have traditionally come in one of two forms — boxed Linux versions, or Windows versions with Linux patches available. As the home Linux market grows, the boxed variety appears to be becoming more popular.

Linux has some way to go before being as ubiquitous a gaming platform as Windows, the variety of titles is increasing.

I SHOULD BE SO LOKI

One of the biggest names you'll see in Linux gaming is Loki Entertainment Software. Loki is responsible for bringing the

majority of big name Linux games to the platform. Loki doesn't write its own games, nor is it a games publisher, though its role involves both working on game code and marketing software. Loki signs agreements with traditional games publishers to bring Windows games to Linux, doing whatever work is necessary to make the game work on Linux, selling special Linux-boxed versions of the game, and providing support. The publisher and Loki then divide the revenue from the title. Though Loki's games are all closed source, it has released a comprehensive set of tools, such as Simple Directmedia Layer and Simple MPEG Player Library, to the open source community.

So far Loki has either helped bring to Linux the following games, or is in the process of bringing them to Linux: Quake III Arena, Unreal Tournament, Tribes II, Soldier of Fortune, Descent III, Civilisation: Call to Power, Myth II, Railroad Tycoon II, Heretic II, Heavy Gear III, and SimCity 3000 Unlimited.

Following are a selection of game demos you can install from the Linux Pocketbook CDs to test your gaming mettle under Linux.

GAME DEMOS ON THE POCKETBOOK CDs

QUAKE III: ARENA

Quake III: Arena, like the popular Unreal Tournament, is a deathmatch-style first person shooter (FPS) involving running around large areas, picking up weaponry and blowing your friends to pieces. If you don't know how much fun these games can be, you need to stay home more.

DESCENT 3

The last incarnation of the classic 3D Descent series is one of the most innovative and popular 3D shooters

The last incarnation of the 3D Descent series is one of the most innovative and popular 3D shooters of all time.

of all time. Find running around large maps on foot shooting your pals a little boring? Then take on a shoot-em-up in three dimensions flying through darkened tunnels blowing up nasty robots intent on causing you and your ship harm. There's a plot in there somewhere, too, but like most FPS games it's little more than a reason to be running around shooting things.

Quake III: Arena — deathmatch gaming at its finest.



playing

RAPTOR

Raptor is a scrolling-style arcade game with an objective to avoid being shot while destroying as many enemy ships as possible. You collect power-ups to equip your Raptor, and at the end of each level you can upgrade your ship with the tokens you have collected along the way. Classic shoot-em-up fun.



TuxRacer is a simple but fun 3D-accelerated racing game.

Linux Games, one of the central Linux game sites.



LINUX GAME SITES

Linux Games

www.linuxgames.com

Happy Penguin

<http://happypenguin.org>

Australian Linux Gamers League

www.luv.asn.au

TUX RACER

Tux Racer is a racing game in which players take the role of Tux, the Linux penguin, as he races down steep, snow-covered mountains collecting herring and avoiding obstacles. The game features full 3D acceleration, and is suitable for game players of all ages. The version of Tux Racer included on the CD is source code, so you'll need to venture into compiling your first game by following the instructions on page 113 in order to play.

IF IT'S ALL TOO HARD

Until Linux becomes more widely recognised as a mainstream gaming platform, it may be easier to simply create a dual-boot Windows/Linux system that allows you to boot Windows for Windows-only games — and boot Linux for everything else. If you've installed Linux on your Windows system, you should already have a working Windows/Linux dual-boot system.

PLAYING WITH OTHERS

The Australian Linux Gamers League is a group of dedicated Linux gaming enthusiasts who provide mailing lists and IRC channels to help newbies use Linux with games, hold events online, LAN parties and more. Check it out for Linux gaming in Australia.



Resources

GETTING HELP 164

KEEPING IN TOUCH 166

UNINSTALLING LINUX 169

COMMAND REFERENCE 170



Getting help

When you're in need of help there are plenty of resources at your disposal, including localised help that you'll find installed on your Linux machine and online help accessible through the Internet.

LOOK TO LINUX

The first kind of help has been discussed throughout this book — the manual files. All commands and most of the important files in Linux are documented with their own individual manuals. Accessing these is as simple as typing:

```
man [command]
```

For more information about this command, see page 125.

HOWTOs

The next source for information is the HOWTO documents. If you look in `/usr/doc/HOWTO` you will find a list of HOWTOs for common tasks, problems and questions. If you're not sure how to do something, this is the first place to look. Each HOWTO is written by a user who has decided to contribute to the community by sharing their knowledge. HOWTOs are also regularly maintained, often by the original author, to make sure that they are always up to date and relevant. More often than not, the HOWTOs contain everything you need to know about the subject, and if you follow the instructions carefully, you're assured of mastering the topic. You will also find a list of mini-HOWTOs in the `/usr/doc/HOWTO/mini` directory. These cover smaller topics than the

AUSTRALIAN LINUX IRC CHANNELS

Server	Channel
irc.linux.org.au 6667	irc.linux.org.au 6667 #linuxaus (Linux Australia) #linpeople (Linux People, 24-hour international support) #humbug
efnet.telstra.net.au 6667	#linux
sydney.oz.org 6667	#linux

normal HOWTOs, but are otherwise similar in nature.

HELP TOOLS

Both Red Hat and Mandrake come with the Gnome and KDE help browsers that allow you to view Gnome and KDE help files, man files, info pages, application help pages and more. Just click on the help button in the Gnome and KDE panels.

LINUX CD

Most of the HOWTOs also come in HTML format for easy browsing. These can be found in `/usr/doc` if you elected to install them during installation. The documentation is also included on the original CDs. Simply browse the CD and follow the links.

THE INTERNET

If you don't find what you need in the documentation that is included, there is literally a world of knowledge awaiting you on the Internet. There are many ways to find information; through Web pages, newsgroups and other users.

IRC

If you've never used Internet Relay Chat (IRC) now is a good time to start. IRC allows thousands of users worldwide to communicate in real time. You will find 'channels' (an area where people can meet and talk) for just about any topic you care to imagine. Linux comes with an IRC client, which you can use to connect to IRC servers and talk with other

LINUX NEWSGROUPS

alt.linux	comp.os.linux.*
alt.os.linux	linux.dev.*
aus.computers.linux	linux.redhat.*

people. Although there are many IRC servers around the world, it's faster and more efficient to connect to local Australian servers. If you're looking for help with Linux, some of the channels you can go to are listed in the table on the previous page.

NEWSGROUPS

Newsgroups allow people to discuss topics in a public forum using message boards — places where messages can be read and posted. As with IRC, there are thousands of newsgroups you can read, however you will need a newsreader in order to do so. Linux comes with a variety of newsreaders (such as knode) as do browsers such as Netscape Navigator. Once you're set up to browse newsgroups with your ISP (contact your ISP if you're not sure how), the table above lists some groups you might visit for help and information.

You might also want to explore Red Hat's mailing list archives at archive.redhat.com. These contain archives of the Red Hat mailing list discussions, which you can use to browse for a particular topic.



Keeping in touch

For most Linux users, the Internet is their lifeline for staying up to date with the system, its applications and its users. Bookmark the following as your staple online diet.

The following list is a compilation of some of the best sites you'll find, what they contain and why you should make a habit of visiting them regularly.

SLASHDOT.ORG

slashdot.org

The catchphrase for Slashdot is 'News for nerds. Stuff that matters.', and indeed, it is an excellent source of information for all things geeky and technologically important. Slashdot has become a well-respected source for news for the Linux community, and is widely read in the IT industry. Visiting Slashdot is the second step in every geek's morning ritual (right after that first sip of coffee).

FRESHMEAT.NET

freshmeat.net

Freshmeat is Slashdot's twin, and focuses on reporting the latest file releases in the Linux world. If you want to see just how active Linux development is, visit Freshmeat and look at the long list of programs that are released every day. If you're looking to see what's new in Linux development, or you're looking for a

particular application, Freshmeat is your first stop.

LINUX.COM

linux.com

Linux.com is rapidly becoming the first place new users learn about Linux. This well-designed, informative site lists plenty of Linux resources as well as Linux news and events. It also has a good tips section for optimising and tweaking Linux. Highly recommended to people who are interested in Linux and want to find out more.

LINUX HEADQUARTERS

www.linuxhq.com

Linux HeadQuarters is an ideal site to monitor the evolution of stable and development kernels. If you're looking for kernel-related resources, Linux HeadQuarters has it all.

LINUX HOWTO

www.linux-howto.com

The Linux HOWTO page is one of the first sites you should visit. If you have any questions concerning how to perform a certain task or any questions about configuring Linux, look through

the HOWTO pages — chances are you'll find an answer here.

GNU PROJECT

www.gnu.org

The GNU Project has been integral to Linux since its inception. It is an excellent site where you can learn about the GNU philosophy and how the GPL (the licence most Linux software is released under) came into being.

ENLIGHTENMENT

www.enlightenment.com

The Enlightenment window manager aims to be both functional and aesthetically pleasing. Packed with options, and still under development, the Enlightenment home page is the first place to look for Enlightenment news and links to the latest version.

GNOME PROJECT

www.gnome.org

Since Gnome is now packaged with both Red Hat 8.0 and Mandrake 9.0, this site needs no introduction. Visit it regularly to keep up to date with the project, and if you can help out, feel free to contact them.

K DESKTOP ENVIRONMENT

www.kde.org

The home page of the most full-featured and powerful desktop available for Linux today. Check it out to see what's new, follow developments and download the latest updates.

PC USER HELP STATION

www.pcuser.com.au/helpstation

Got a Linux question? Looking for an answer? The Linux section of PC User Help Station might just have what you're looking for.

RPMFIND.NET

rpmfind.net

If you're looking for a particular program or type of program, Rpmfind.Net hosts the largest repository of Linux packages in the world. Simple to navigate and search, Rpmfind.Net is an excellent resource to find packages for any RPM-based distribution.

LINUX WEB WATCHER

www.webwatcher.org

Linux Web Watcher keeps track of hundreds of Linux sites, logging sites that have recently been added or updated. It's also an excellent place for searching for Web pages on particular topics, as the sites are indexed by category.

THEMES.ORG

themes.org

Themes.org is the only place you need to visit to spice up your desktop. Hundreds of themes are catalogued here, covering window managers such as Enlightenment and Window Maker as well as operating environments such as Gnome and KDE. Once you've set up your Linux box to your liking, themes.org is the next step in creating a totally customised desktop environment.

LINUXBERG

www.linuxberg.com

Linuxberg is the Tucows for Linux.



The resource is well designed, easy to navigate and fast to download. Although Freshmeat covers a much wider scope when it comes to new programs, Linuxberg is a great resource for the new user, as programs are given a rating and are broken down by category. If you're a fan of the Tucows Windows archives, you'll feel at home with Linuxberg.

Lastly, one FTP site every Linux user should bookmark:

AUSTRALIAN KERNEL ARCHIVE

ftp.au.kernel.org

When you want to download the latest stable or development kernel, the Australian kernel archive is a mirror of the US kernel archives and is the easiest and fastest method to grab a new kernel.

For an extensive list of Linux sites, check out the Pocketbook CDs. If you're connected to the Internet at the time, you can just click on a link to visit a site.

USER GROUPS

There is at least one Linux user group in each mainland state of Australia.

Most groups have a discussion email list to which you can subscribe, regular meetings (often monthly) you can attend and a Web site for information about the group.

AUSTRALIAN LINUX USER GROUPS

CLUG	Canberra Linux Users Group	samba.anu.edu.au/clug
HUMBUG	Home Unix Machine, Brisbane Users Group	www.humbug.org.au
LinuxSA	Linux Users Group of South Australia	www.linuxsa.org.au
LUV	Linux Users of Victoria	www.luv.asn.au
MLUG	Melbourne Linux Usergroup	www.mlug.org.au
PLUG	Perth Linux Users Group	plug.linux.org.au
SLUG	Sydney Linux Users Group	www.slug.org.au
USQLUG	University of Southern Queensland Linux Users Group	www.sci.usq.edu.au/lug

Uninstalling Linux

Linux isn't for everybody. If you want to remove it from your system, uninstalling Linux is almost as easy as installing it.

To remove Linux from your system and LILO as your boot manager, follow the steps outlined below:

REMOVING LILO

This is the easy part. Just boot your Windows partition, open a DOS prompt and type in the following:

`FDISK /MBR`

You can also do this from a DOS boot disk. This command will restore the Master Boot Record, overwriting LILO. If you use third-party boot manager software, you can remove Linux as an option from the list.

REMOVING LINUX PARTITIONS

Because DOS and Windows are unable to recognise Linux partitions, their partitioning software, FDISK, is unable to remove them. If you don't have commercial partition management software such as PartitionMagic, you'll need to use the Linux fdisk program.

You can simply start a Linux installation and use the partitioning software to delete the partitions, and then back out of the installation,

or alternatively boot your Linux boot disk and run fdisk. If you choose to do the latter, you'll find that fdisk doesn't operate like its DOS counterpart.

To delete partitions with fdisk first make sure you know which partition numbers you want to delete. Press P and look at the listing. Then, press D to delete a partition and type in the partition number. Remember to delete both the Linux and the swapfile partitions. When that's done, write the changes with W and reboot your machine.

In order to take advantage of the space which was previously occupied by Linux you can either make your existing Windows partition larger using software such as PartitionMagic, or simply create a new partition in its place, which you can then use as an extra drive for Windows or any other operating system. To do this, use the DOS FDISK program to create a new partition, and then be sure to format it with the correct filesystem.

In the unlikely event that you have trouble removing Linux, see that start of this chapter for help.



Command reference

Where would a Linux geek be without a quick and handy command reference to help manage and maintain their sexy new Linux box?

This command reference is by no means a complete guide to commands under Linux. It will, however, provide you with all the information you need to perform basic system management and maintenance tasks.

We've included a comparative listing of DOS and Windows commands, as well as basic commands for user management and permissions, process management, installing software and everyday tasks. Don't forget to use the man and info commands to find out more about how and where these commands are used.

DOS/WINDOWS AND LINUX COMPARED		
DOS COMMAND	LINUX COMMAND	PURPOSE
CD	cd	Change directory
DIR	ls	List the contents of a directory
CHDIR	pwd	Shows your full location in the filesystem
MOVE	mv	Move files
COPY	cp	Copy files
DEL	rm	Remove [delete]
DELTREE	rm -r	Recursive delete [wipe a directory and its contents]
REN	mv	Rename files
EDIT	pico	Edits files with a simple text editor
TYPE	less or more	View files

USER MANAGEMENT AND PERMISSION COMMANDS		
COMMAND	FUNCTION	BASIC SYNTAX
su	Login as another user account	su - [user]
adduser	user account (only available to root)	adduser [account name]
passwd	Change a user's password	passwd [account]
chmod	Change permission for files and folders	chmod [permissions] [target]
chown	Change ownership of files and folders	chown [user]:[group] [target]

PROCESS MANAGEMENT	
COMMAND	FUNCTION
ps	Display processes currently running
ps aux	More informative ps display
top	Automatically refreshing list of running processes
kill [process id]	Kill a process with the specified process id
kill -9 [process id]	Kill a stubborn process

INSTALLING SOFTWARE	
COMMAND	FUNCTION
rpm -ivh [package file]	Install a new package
rpm -uvh [package file]	Upgrade a package
rpm -e [installed package]	Uninstall a package
rpm -q [package]	Find out if a particular package is installed
rpm -q --requires [package file]	Find out a package's dependencies
rpm -qlp [package]	List all the files stored within the package
rpm -qix [package]	Read a description of an installed package
rpm -qixp [package file]	Read a description of an uninstalled package
rpm -v [installed package]	Verify an installed package

COMMANDS FOR COMMON TASKS		
COMMAND	FUNCTION	SYNTAX
locate	Locates a specific file	locate [filename]
ln	Creates a link to another file or directory	ln -s [target] [link name]
xhost	Add/remove X access to your display	xhost +[host name]
du	Shows disk usage of a file or directory	du [directory]
df	Shows free disk space on mounted filesystems	df
tar	Create and extracts the contents of .tar and .tar.gz files	tar -zxvf [filename.tar.gz] Extracts archive tar -cxvf [source] [filename.tar.gz] Creates archive
tail	Shows the bottom lines of a text file. Generally used with log files to show the most recent data	tail [options] [filename]
reboot	Reboots the machine	reboot
shutdown	Shuts down the machine	shutdown [options] [time]

POPULAR CONFIGURATION FILES

APPLICATION OR SERVICE	CONFIGURATION FILE	DESCRIPTION	INITIATE CHANGES
Samba	/etc/samba/smb.conf [RedHat] /etc/smb.conf [Mandrake]	Configures Windows file and print sharing using the SMB protocol. Samba allows Linux and Windows machines to share files, printers, account information and more.	Service samba restart
LILO	/etc/lilo.conf	The Linux Loader. LILO is a small program that launches when you start your machine, and allows you to choose between different operating systems, as well as different Linux kernels.	lilo
Filesystems	/etc/fstab	The filesystem table stores the details for all your automatically mounted filesystems, as well as default options for temporary filesystems.	No action needed, unless non-Linux filesystem permissions were modified, when a user should log out and back in again
Modules	/etc/modules.conf	Specifies which modules to load at boot time.	Changes will be seen on reboot
X	/etc/X11/XF86Config	Specifies your graphics card and monitor, display size, display resolution and colour depth.	Log out and restart X
Default runlevel	/etc/inittab	There are six different runlevels on Linux. Each 'runlevel' changes the services that are running. For example, runlevel 5 has a GUI login, runlevel 3 is console based, and runlevel 1 has only a small set of services, and is used for troubleshooting. Edit the initdefault line to specify which runlevel to boot into.	Changes will be seen on reboot