NEW

LINUX FORMAT

The UK's ONLY dedicated
Linux magazine

www.linuxformat.co.uk

**Everything you
need to install &
run the Linux OS**

Join
the
revolution!

Windows
buster

25 pages of Tutorials,
How-tos and tips

Everything from
Apache to Zip drives
explained this issue!

Apple Networking

Fool Macs into thinking
your Linux box is
one of them!

Open Source?

What does it all mean,
and is it a Good Thing?

Which distro?

All the popular Linux
flavours rounded up
and rated.

DEFINITE LINUX 7.0

**Full Linux Operating
System on this CD!**

Explore the exciting
world of Linux with
this fully-featured
UK-based distribution
of the world's best
operating system!

PRESENTED WITH ISSUE ONE OF LINUX FORMAT

For details of supported hardware, please check documentation on CD.

BEGINNERS START HERE!

Full Linux Distribution on the CD,
plus an exhaustive ten page guide
to installing your new OS.

LXF01 £4.99 (£6 outside UK and Eire)
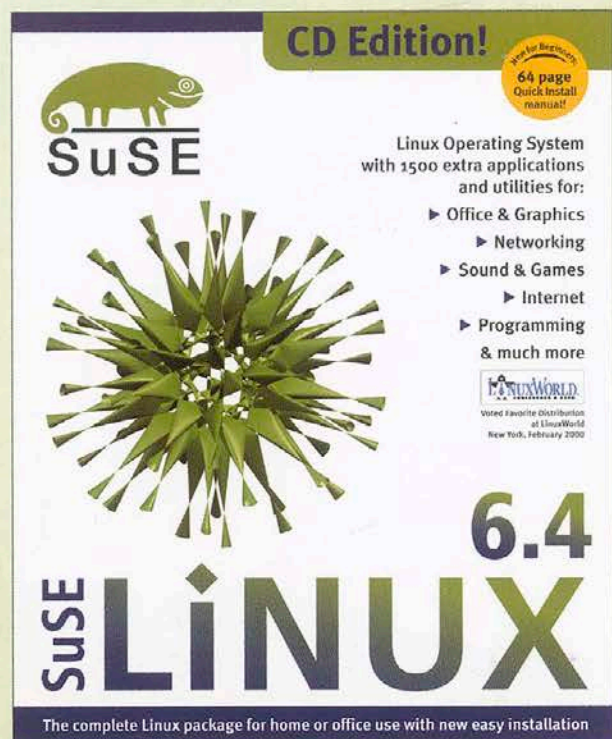
future
Media With
Passion

MAY 2000   £4.99

0 5

9 771470 423002

WELCOME

You hold in your hands the very first issue of the very first UK newsstand magazine dedicated exclusively to the Linux operating system. The Linux community has been without such a magazine for too long, so we're here to put things right!

As you can see from our aims listed here, we exist not only for current Linux users, but also to help welcome converts into the fold. To this end we have a special beginners' guide in this issue, starting on page 48, and a full Linux distribution on the CD. I hope veteran Linux users won't be too upset about the space this takes up – we all have to make a start somewhere, after all!

The mainstay of this magazine will be its features, which we hope will entertain and inspire you, the reviews and roundups to guide you through the ever-growing range of Linux software and our tutorial section, which should provide a handy guide for Linux advocates of all degrees of experience.

We want to create the best magazine we can for all our readers, and all serious suggestions will be considered, so tell us what you'd like to see and we'll do our best to give it to you! There are plenty of contact details here so you can easily get in touch with the relevant people. Go on, don't be shy – we all know Linux users are a vocal bunch, so let's put your opinions to some good use!

Nick

Nick Veitch
Editor
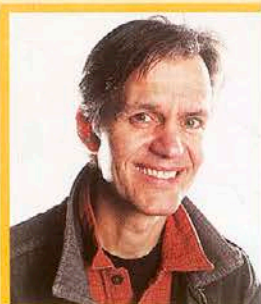
Chris Crookes
Art Editor
He does the design and pictures.
Don't worry, he can't hurt you.
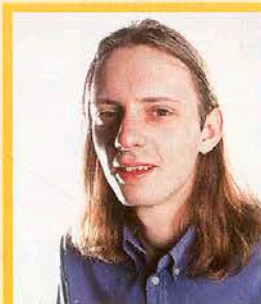
Richard Drummond
Staff Writer
His words of wisdom are
scattered liberally within.

The aims of this magazine...

Linux Format is a magazine dedicated to Linux and the Open Source community. The aims of this magazine are quite simple:

● To help our readers get more out of their Linux experience, through our tutorials, features and advice pages.

● To provide Linux Users with accurate and unbiased information.

● To promote Linux and the Open-Source movement.

Contact us

Letters for publication:
lxf.letters@futurenet.co.uk

Subscriptions/back issues:
subs@futurenet.co.uk

Technical help/Ask the Experts:
lxf.answers@futurenet.co.uk

Disc problems:
lxf.support@futurenet.co.uk

General enquiries:
linuxformat@futurenet.co.uk

Website
www.linuxformat.co.uk

Or send your letters to:
LINUX Format, Future Publishing,
30 Monmouth Street, Bath, BA1 2BW
Phone: 01225 442244 Fax: 01225 732295

CONTENTS

Welcome to the first ever issue of **Linux Format**, your guide to all things Linux!

Linux in Action
Linux scoops an award for enterprising scientists_ p14

Three young scientists and their trusty Linux box scoop major science award! Read their story here.
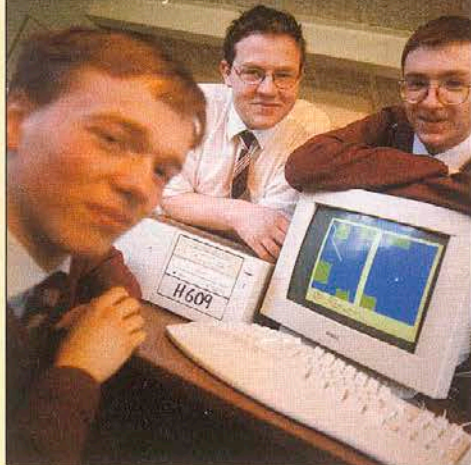
Licensed to thrill?
Open Souce licencing. What does it all mean, and is it a good thing? _____ p20

The intricacies of licencing explained. What's the difference between the licences and why are there so many?
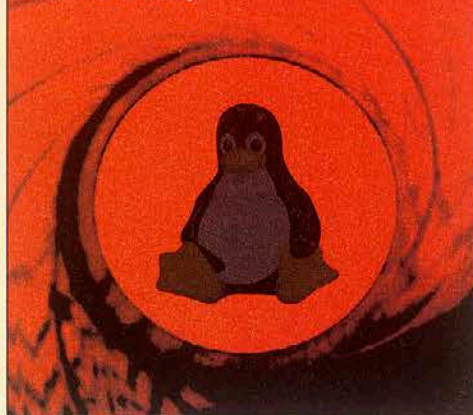
Apple Networking
Fool Macs into thinking your Linux box is one of them! _____ p67

Connecting your Linux box to an Apple Mac network couldn't be easier - if you've read our easy to follow tutorial that is...

Everything you need to install & run the Linux OS

Join the revolution!

Thinking of taking your first steps with Linux? Don't be afraid, we're here to help you discover this exciting new way to enhance your computing experience! Our Newbie's guide will tell you all you need to know _____ **p48**

LINUX Format CD

DEFINITE LINUX 7.0

Full Linux Operating System on this CD!

Explore the exciting world of Linux with this fully-featured UK-based distribution of the world's best operating system!

PRESENTED WITH ISSUE ONE OF LINUX FORMAT

For details of supported hardware, please check documentation on CD
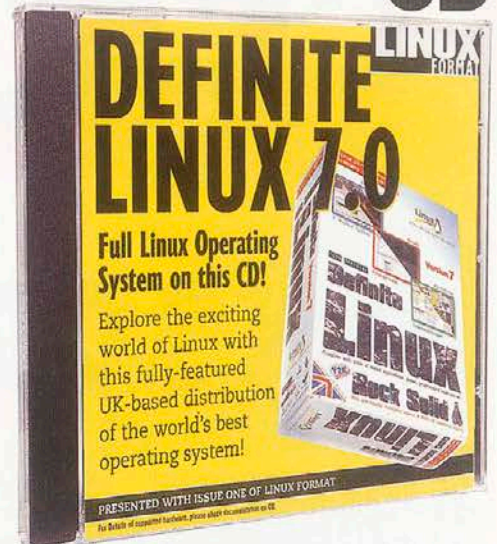
FULL LINUX OS ON OUR CD THIS ISSUE!

There is a complete version of Definite Linux 7.0 on our CD this issue, all ready to be installed! This is a pretty up to date distribution based on Red Hat, suitable for beginners and experts alike!

And

Have your say about what you'd like to see on our CD – take part in our survey! Details are on page 84

LINUX FORMAT on the CD

Wherever you see this logo it means there's related stuff on the CD

Please read the coverdisc instructions on page 84 before installing from the CD!

NEWS

New browsers, new PPC support, new commercial software ports, new websites, new server software, new UK offices, new games consoles? – it's all in the **Linux Format** news pages!

Netscape 6 is nigh

A preview release of Netscape's next generation browser is now available for download from their website. *Netscape 6 PR1* is the first public Netscape product to make use of the new *Gecko* engine from the open source *Mozilla* project. It is claimed to be smaller faster and more standards compliant than previous versions.

Netscape seem to have borrowed more than just Gecko, though. Users who have tried out the latest beta release of the *Mozilla* browser, dubbed *Milestone 14*, will be hard-pushed to notice any difference with *Netscape 6*. The 'skinnable' interface sports an identical look, the Netscape 'N' transfer animation being the only obvious difference.

With release 6, Netscape have dropped the *Communicator*

distinction, but it contains a similar set of integrated components to before: *Navigator* (the browser itself), *Netscape Mail* and *Composer* (the HTML editor). New features include *My Sidebar* and the Instant Messenger, a client for the AOL instant messaging service. *My Sidebar* appeared in *Mozilla* too and is an area of the browser window which the user can customise. It can display search results, bookmarks and 'What's related' links, plus a range of embeddable applications. There's also a new cookie browser and Personal Security Manger.

The *Netscape 6* preview is available for Windows, MacOS and Linux. The Linux version requires a 2.2 series kernel, glibc2.1, and 48MB of RAM, and it weighs in as a 10MB download. You can go to:
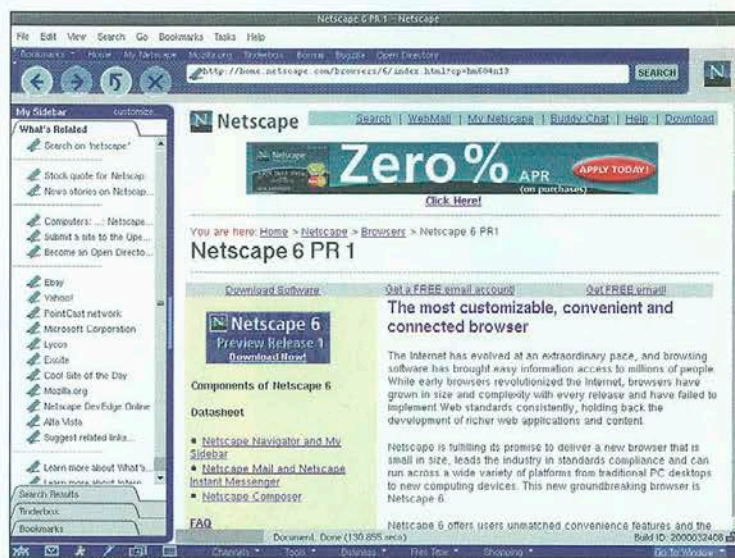
The shape of browsers to come? You can check out Netscape 6 now!

http://www.netscape.com/browser s/6/index.html to get your copy.

Remember that this is beta software; it has lots of rough edges to polish off.

SuSE now for PPC

Version 6.4 of SuSE Linux is now being released for all the usual suspects, but a PPC version will be following close behind.

SuSE has a reputation for good, solid and reliable Linux distributions, so it will be interesting to see how it copes with the varyingly capable G3 and G4 Macs.

The PPC version is the product of some six months of hard work by the

SuSE team. The first beta was shown at the Mac World show in January, and went down pretty well.

A SuSE insider told us the progression to PPC wasn't that hard: "SuSE uses a sophisticated automatic build procedure for the software packages contained in the SuSE Linux distribution. In conjunction with the necessary modifications of the Linux kernel done by the Linux community

and members of the SuSE Linux Labs, this enables us to reliably produce versions for other hardware platforms, with all the secure benefits of the auto-build system."

The Mac platform is currently attracting a lot more interest from Linux users due to the particularly good performance/price ratio on the G4, as well as Apple's current interest in all things open source.
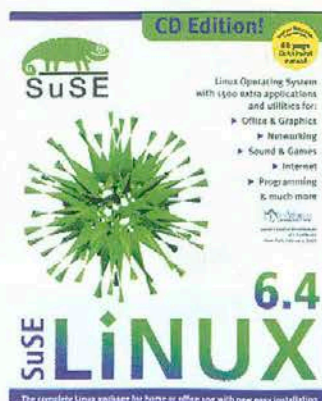
Apple recently released Darwin, the prototype for the forthcoming OS X, as open-source software.

COREL GET "AGGRESSIVE" with Linux

Canadian software giants Corel, best known for their once untouchable PC drawing package, *Corel DRAW*, have announced a raft of new release dates for Linux versions of some of their most popular products.

Corel have already ported

CorelDRAW9 will also bring with it easy access to a huge range of graphic formats.

These are Mac screenshots, the Linux version will be similar, but different. Probably

Photopaint, Corel DRAW and now Ventura will be available on Linux systems soon!

WordPerfect, and now *WordPerfect Office* to the Linux platform, and have also announced their intention to port *Corel DRAW* and *PHOTOPAINT*. Now it seems that these will be joined by a simultaneous release of *Ventura Publisher 8.5* on Windows and Linux platforms.

Ventura has been around for quite some time on the Windows platform, and although it has lost out in the DTP stakes to big guns Adobe and Quark, the software is nevertheless quite capable. The new version of *Ventura* will feature enhanced PDF and web publishing capabilities. "Our inventive and aggressive Linux and Windows initiatives have enabled us to develop simultaneous versions of *Corel VENTURA Publisher 8.5* for Linux and Windows," claimed Corel VP Derek Burney.

Corel launched their own version of Linux in November last year, attracting just as much acclaim as it did flames on the newsgroups — when it works it is very very good, but

when it doesn't... well, see our roundup starting on page 38 of this issue.

The last 12 months have seen a re-invention of Corel. Although still behind their flagship title *DRAW* in the PC market, a great deal of diversification has been going on - mainly into the Linux market. In fact, there are more Linux launches on the schedule at the moment than Windows ones.

What remains to be seen is what impact the appearance of 'famous' commercial software will have on the Linux community. On a features basis, *PHOTOPAINT* doesn't offer much over the *GIMP* software, but feature-wise, *Corel DRAW* is way ahead of any open-source drawing software for Linux.

This might not be the end of high-profile software appearing form the Corel stable. In a surprise move, they have also acquired the rights to several former Metacreations products, including the landscape/3D rendering software *Bryce* and the popular natural media art package *Painter*. Corel have
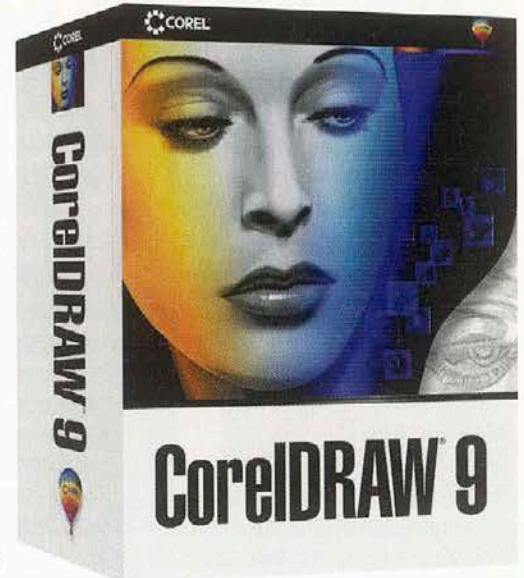
CorelDRAW will be shipping earlier than expected. Will it soon be joined by *Painter* too?
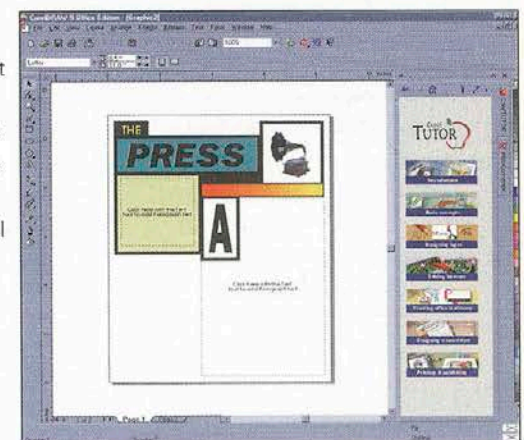
CorelDRAW was once king of all the structured art packages by a long way.

made no indication that future versions will be released for Linux, but since the rest of their major software products will have made the leap before the end of the year, we'd bet our money on there being some more interesting announcements soon. For more information, visit: **www.corel.com**
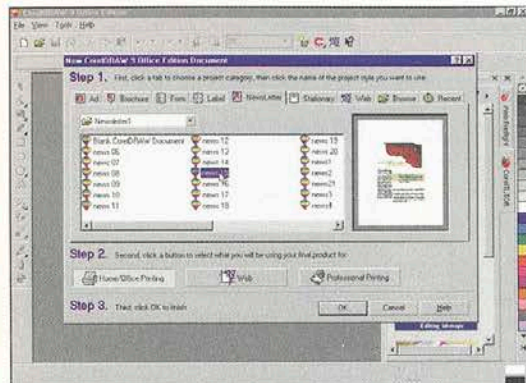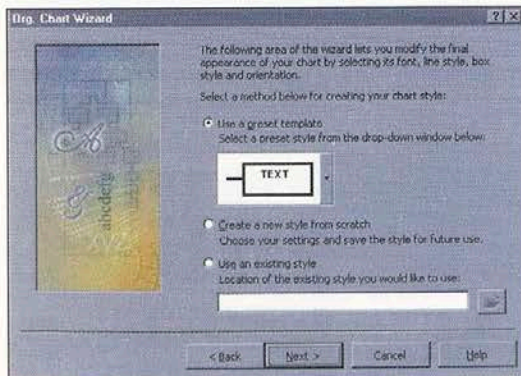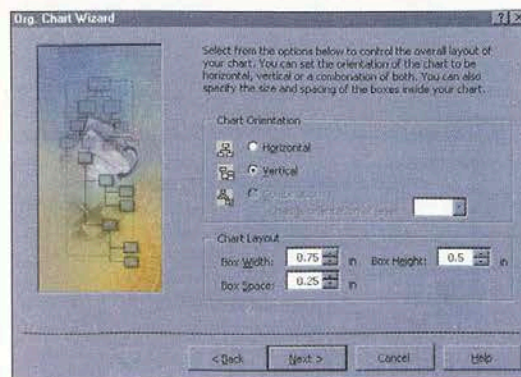
THE RELEASE SCHEDULE

WordPerfect Office 2000	Out now
Corel PHOTOPAINT 9	June
Corel DRAW 9	July
Corel VENTURA Publisher 8.5	Later this year

Aladdin Expander for Linux

Aladdin Systems, best known for their Mac software range, including *Stuffit Expander* and *Anarchie*, have released a beta version of *Expander* for Linux. Expander de-archives compressed and encoded files and has support for many archive types on all platforms, as well as its own, native ".sit" files.

Other types supported include zip, gzip and Arc, as well as encoding schemes like MIME, BinHex, MacBinary and uuencode. The beta version of *Expander* has been designed for a 586 running Corel or Red Hat Linux and is available for download from the Aladdin site at: **http://www.aladdinsys.com/expander/**

COBALT ACQUIRE CHILI!SOFT, AND PRODUCE NEW ASP-CAPABLE SERVERS.

Those of you lucky enough to own some of their hardware will know that Cobalt produce a range of powerful, compact and generally cool servers for office networking.

Well, they have some more good news for people running webservers for Internet or Intranet use. Cobalt have acquired Chilli!Soft and their ASP software for Linux, which can now be included in any of their RaQ 3 servers. For more information, visit: **www.cobalt.com**

TURBOLINUX IN THE UK

TurboLinux, one of the more popular Linux distributions, have announced the opening of their UK offices for support and distribution. A new version of TurboLinux is due pretty soon. You can check out the state of play at **www.turbolinux.com**. In the meantime, you can call their UK office on **01752 313190.**

GIMP NOW ONLINE!

A new website has launched which be of interest to any Linux user who does a lot of image manipulation. Spencer Kimball, one of the creators of Linux's finest graphics tools, the *GIMP*, has started a web-based photo manipulation service.

Users are given over 50MB of web space in which to upload their raw photographs. From there, the images can be processed with a variety of effects, shared with others in a library, compiled into an album or even sent as e-postcards.

The power behind the website is 10 Linux boxes running a special version of *GIMP*. The Linux machines provide all the power to apply the photo effects.

"There are many websites today that help consumers display their photographs on the web and print them out," said James Mattis, Vice President of Business Development for Online PhotoLab, "but until now there has been no way for web users to get the crucial digital enhancement tools that are so common in the photography world. With Online PhotoLab, people can get all the power and flexibility of great graphics
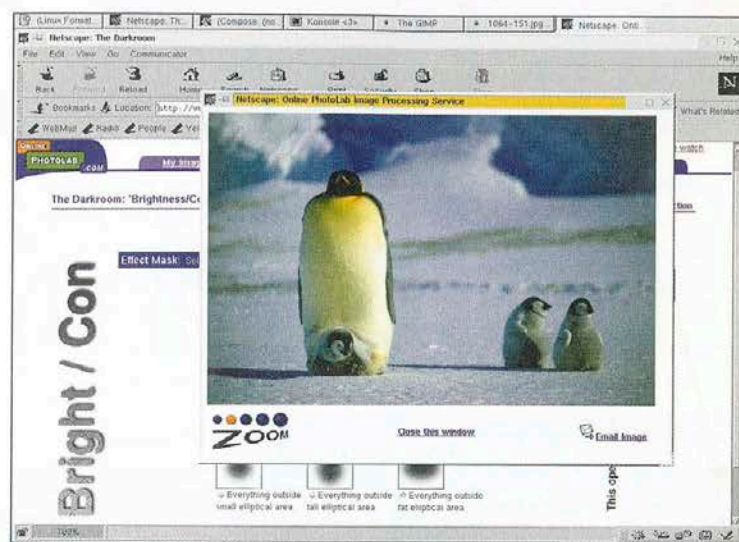
Aren't they cute... right, now lets mess about with them with the *Gimp* **powered transforms.**

An online thumbnail gallery lets you view your uploaded (and transformed) images.

software for free. Combining the robustness of the Linux operating system and cost effectiveness of commodity hardware, Online PhotoLab has been designed to provide advanced functionality under extreme load with extremely low Cost Per User."
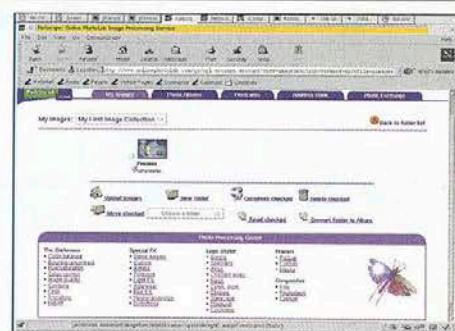
But why not try it for yourself? After a simple registration process, you can store large numbers of images online, swap them with others and more. You'll find the Online Photo Lab at this address: **www.onlinephotolab.com**

One of the neat features is a postcard design service.
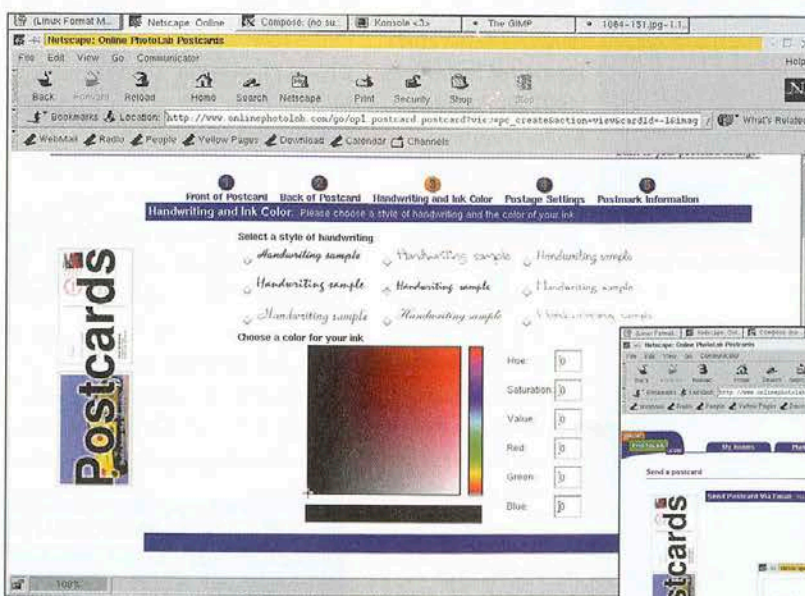
As well as transforming images, you can also use them to create custom postcards (above and right) or albums,
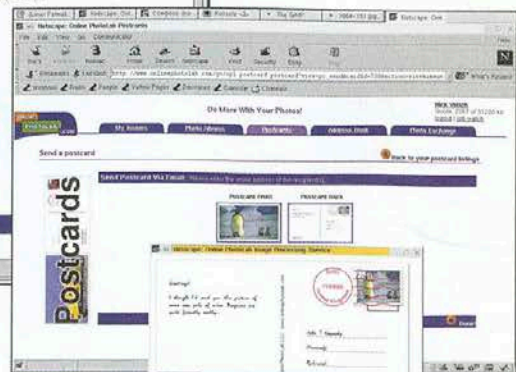
LINUX CONSOLE is GO!

Sony executives may not be losing any sleep yet, but a Linux console is definitely on the way. Uninspiringly titled the L600, the console is the brainchild of Indrema, a technology startup in the US.

The L600 will be based around a 600MHz processor, and it's planned to have a broad range of network connectivity. The plan is not just to develop a Linux-compatible games console, but to create a hybrid computer, games, Internet and interactive TV device. Designed very much for set-top use, the system will incorporate MP3 recording and playback facilities and will have an optional infra red keyboard. Needless to say, exact specifications aren't yet forthcoming, and neither is a price.

However, this Linux project does seem to have more substance than many of the putative Linux devices clogging the Internet ether, and is likely to be shipping this year, according to some industry insiders.

Your home
will never
be the same.
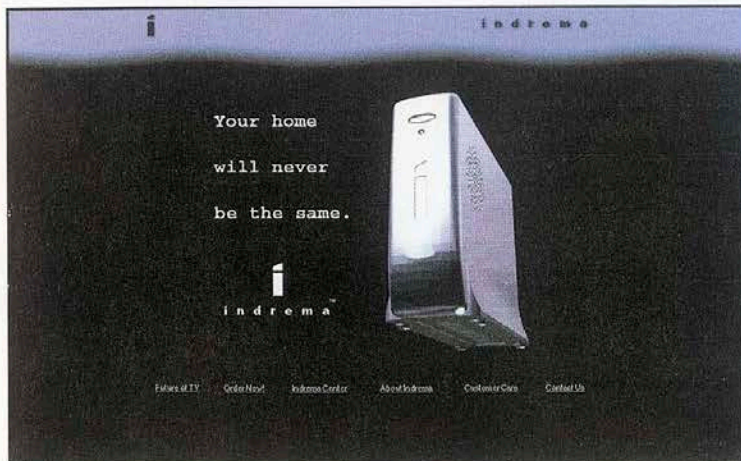
It would be a lie to say there is a wealth of information on the website, but we expect there'll be more details there soon.

SOURCEFORCE finds Trove

SourceForge, provider of free hosting services to open-source developers, are to begin implementing the Trove system for categorising projects hosted on their site. As a first step, SourceForge are asking project maintainers to classify their projects according to the Trove categories.

Trove is still in its infancy, but is planned to be a database system for managing software archives. It will provide an interface that will ease the job of the archive maintainer and make it easier for users to locate the software they need. The initial Trove design was created by open-source advocate Eric Raymond, and can be found at the following site: **http://www.tuxedo.org/~esr/trove**.

SourceForge host many high-profile open-source projects, including Aladdin *Ghostscript* and the *SawMill* window manager. For more details, see: **http://sourceforge.net/**.

COMMENT

Does Easy mean Bad?

When I started work on creating Linux Format, I had a conversation with a contributor I didn't know very well (I'll save his blushes by not naming him) and he asked me what version of Linux I was running. "Mandrake 7.0," I replied, and was met by a long pause. "Yes, that's right, the really simple one that's for thickies," I added, just to make his point for him.

I was quite surprised that this sort of prejudice still went on. It is, I suppose, only to be expected in a community which has its roots in a time where real men (very few women were allowed) wrote their own applications, drank pure caffeine and had object oriented dreams.

But to label this behaviour as prejudice is, these days, implying that it's wrong. Is it wrong? Is Mandrake just for "thickies", who aren't going to be using Linux "properly"?

I ended up with Mandrake on my Linux box at work, mainly because it was the last in a series of current distributions I installed while bringing myself "up to speed" with the current versions. I've almost got it set up how I want it now, so it's likely to stay for some time.
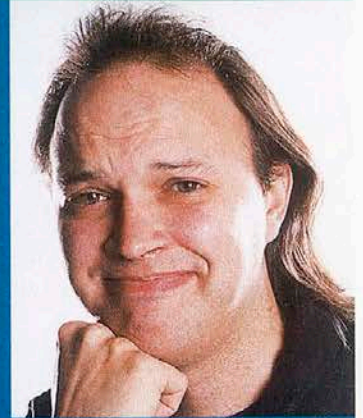
The graphical installer might seem a bit over the top, but it works just the same. You don't have to keep KDE as your default window manager. You don't have to use the lothar configuration tools to set up your sound card. You can change it in pretty much the same way as any other Red Hat-based distribution, so why should there be a stigma attached to it being slightly more user-friendly? I'm all for people understanding about their computers, but how are most people going to learn if they are frightened of the install process?

If I had told my contributor friend that I previously ran NetBSD on a PPC Amiga, which took about two weeks to configure and required a lot of hacking of X servers to run properly, no doubt I would have risen in his estimation, but to tell you the truth, if I could have installed it in half an hour with a nice front end, I certainly would have done it that way instead...

DO YOU HAVE A HOT NEWS STORY?
The best way to let the Linux world know about it is to have it appear here! Email us your news at linuxformat.co.uk

SICK OF CRASHING?
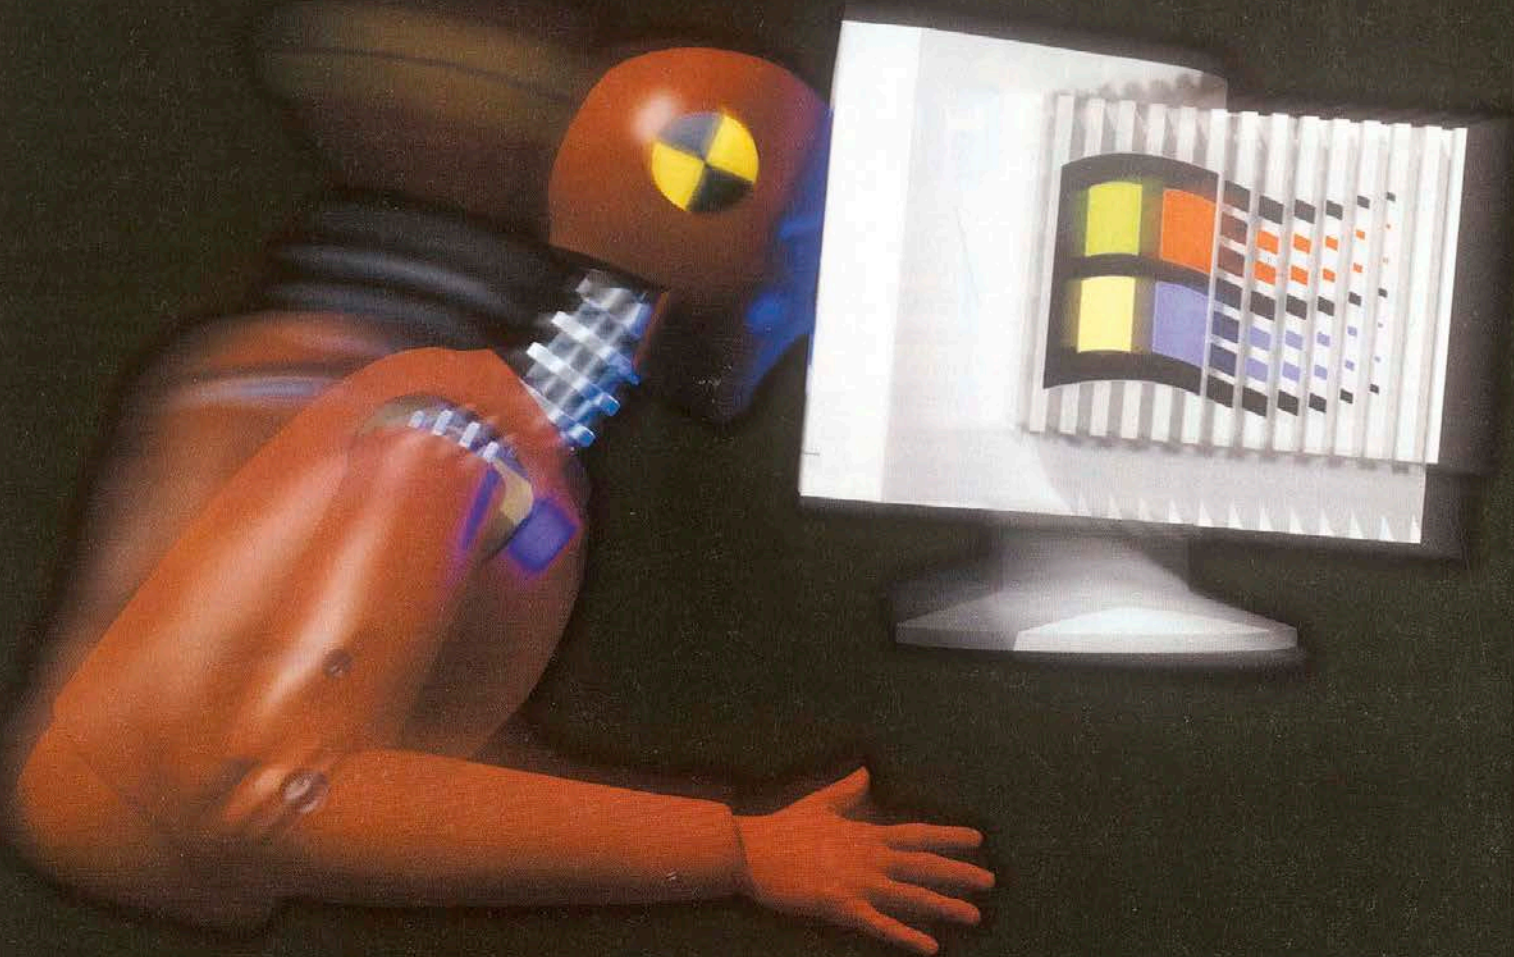
SWITCH TO

PENGUIN COMPUTING™

The World's Most Reliable Linux Systems

"Penguin claims to have 'the world's most reliable Linux systems.' A look at the system suggests they take that claim seriously." -Linux Weekly News

"Penguin Computing's systems, which are designed to run Linux, were the only ones on which we had no installation problems." -PC Week

SCIENCE IN

H609

ACTION

Linux trio win award

College students Darius McGann and Trevor Johnston discuss their prize-winning project studying artificial intelligence which evolves and learns through a process of trial and error...

Linux, the ultimate alternative OS, this year played a large part in an award-winning project entitled, 'AI: Future or Fable?' at Ireland's Esat Telecom Young Scientist and Technology Exhibition in the RDS, Dublin.

The project was designed to demonstrate the possibilities of computer learning and consisted of a simple, yet effective, program coded in C++ and displayed through the unusual medium of *Blitz Basic*, all run on a Linux machine. Entirely our own creation, the project won the Nortel Networks Award for Best Use of Information Technology.

Three students from Saint Oliver's Community College, Drogheda, Co. Louth, entered: Trevor Johnston (myself), Alex McAllister, both of us sixth (final) year pupils, and fifth year →
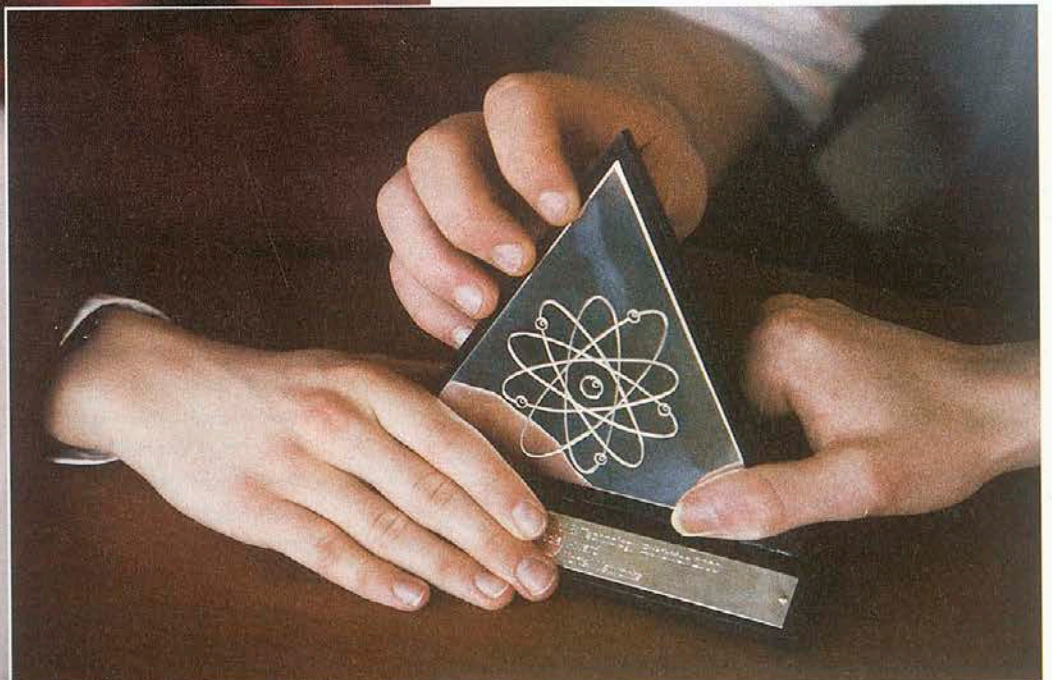
The award itself, which seems to have strange powers of magnetism.
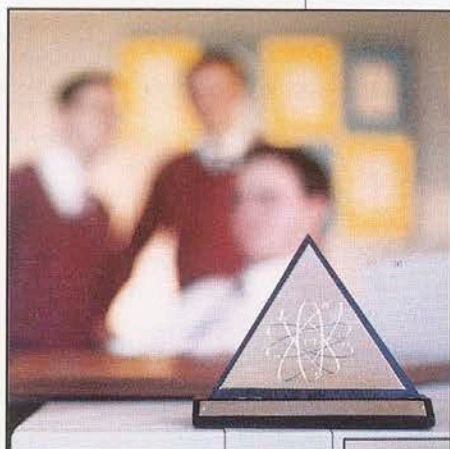
Hmmm... Maybe if they keep their razor-blades under it they'll stay sharp.

→ group leader Darius McGann. We didn't know what we were getting into when we decided on the topic. As we had left ourselves in the unenviable position of posting our entry forms late, Darius was forced into making a quick decision on the subject matter. However, his choice of artificial intelligence (AI) was also of interest to Alex and I, and provided plenty of scope for an interesting display.

Considering that a general form of AI has not yet been discovered, we were in trouble regarding actual material work. Our only solution was to investigate if man-made intelligence might be possible. After several long debates on exactly how to accomplish this, we decided on our own implementation of one aspect of AI, namely computer learning. This, at least, could be achieved through some hard work and didn't require feats of mastermind ingenuity.

We were adamant that the project was going to be all our own work. Apart from the obvious happy side-effect of saving us time getting to grips with 'real' AI languages like Lisp and Prolog, or learning the ins and outs of neural network theory, the idea would, in any case, be a far more worthwhile endeavour if we had to develop everything ourselves. Our program learns by trial and error the fastest way of completing its given task; two coloured

blocks move around a series of randomly generated two-dimensional maps. White moves randomly, while black has two different patterns of movement. Black analyses the average times for him to catch white and develops a preference for the pattern that gives him the fastest results. In mathematical terms, the likelihood of a given pattern being chosen is inversely proportional to the average time taken by that method. An element of randomness compensates for any misleading test runs. In this way, the program can be said to have learned by trial and error – similar to the way in which humans learn.

THE POWER OF LINUX

The program, the unfortunately titled 'Virop', was written entirely in 'super-mega-compressed'-time, recording its activities to a series of text on a Linux machine, using the GNU C++ compiler to run the algorithm and an Amiga emulator to graphically display the algorithm's work. The C++ portion plays out a whole set of games in what could be called files, which just happen to lie in a directory that will later be used by the emulated Amiga program. That half of the program will then 'play back' the games for the benefit of human observers. Why the choice of Linux and the more exotic choice of the Amiga? For Linux, there were many compelling reasons. Quite apart from the fact that, as our resident coder, I am a full-time Linux user, we just would not have tolerated something like Windows for the production of the project. Suppose we had barged ahead with Windows. Singly most important was the C++ compiler. Install any Linux distribution and you are nearly guaranteed to have a

C and C++ compiler there and then. A freshly installed Windows box usually requires a couple of hours of work to find, download and correctly configure a free C++ kit. Now, when you get over the

Linux proved to be the tool they needed to develop their award winning project.

initial pride of getting this to work, you will probably need a text editor. Linux users have a large choice, among them *Vi* or *Emacs*. Both have extensions for C programming, but those who chose to install Windows have to choose between *Notepad* and *WordPad*. These two ever-so-slightly more primitive editors (putting it mildly) will also require some trial and error learning of your own to sort out the filename extensions of your source code.

Perhaps Windows wins with respect to the Amiga emulator. Although UAE is developed on Linux, the Windows version is probably slightly easier to install. However, this is basically a moot point as any competent Unix user will have no trouble getting it working.

Trevor reckons the pretty display on screen had as much to do with the award as the science...

Foregoing the luxury of a decent music player, you now have a basic development setup. Once you get accustomed to the confines of a single desktop, single user environment and a single choice of user interface, you can start coding away: just pray that your operating system doesn't seize up too much when you start pushing it to its limits. Which will drive you mad. Forgive our strong Windows criticism, but the above is just as applicable to every other popular non-Unix operating system today. I think all this just goes to show why Linux was the sensible choice in our case. In no way does Windows prove more suitable. The Amiga emulator, UAE, is a processor-intensive application of near-biblical proportions. Linux had absolutely no problems keeping up with this, five copies of the *Emacs* text editor and as many command prompts as we needed, running all at the same time – try doing that on the same hardware running Windows. That we could also simultaneously run an MP3 player while exhibiting for long hours in the RDS made us even more indebted to Linux.

We had a nice machine for coding on, but nothing spectacular. A 333Mhz AMD system with 64Mb of RAM, the X window system never strained under the load of running all this simultaneously, a great testimony to the strength and reliability of free software.

WHY CHOOSE THE AMIGA?

Now, why Amiga? A popular home computer in the late 1980s and early 1990s, it had graphics and sound far ahead of its time. What really made it unique, though, was its 32-bit, windowed, multitasking operating system. This was in 1986. *Blitz Basic* was (and is) a language based on BASIC, but which also had full support for the Amiga's custom hardware. Armed with this programming tool, the platform still offers the advantage of easily implemented graphics. *Blitz* was, for a while, the language I used and I insisted it was the only way we were going to produce a nice display in a hurry. It proved ideal for the eye-pleasing, smooth, split-screen scrolling effect we finished with. Emulation was necessary to link the C++ program and the graphical front-end, and Linux made this easily possible. The Amiga may be ageing, but it's obviously got some life in it yet, at least as far as the OS goes.

Now all we had to do was travel to the RDS and showcase our work. The Young Scientist Exhibition is not a typical science fair. It runs for five days and is open to the public for much of this. Consequently, the pretty display on our Linux box was no doubt a factor in our eventual capturing of an award, appealing as it did to the exhibitionist nature of the event. We learned how beneficial Linux can be to scientific use and, as people have since pointed out, only a few years ago it would have been inconceivable to have an operating system of Linux's power for no cost. In

A few years ago, it would have been inconceivable to have an OS of Linux's power for no cost...

an educational context, its flexibility has allowed our school, situated 30 miles north of Dublin, to give email addresses, web space and file space to each and every student, making it the most wired school in Ireland. This system has been running there for over a year now with barely a glitch, and has benefited school life enormously. Linux games and office software may be a bit thin on the ground, but the quality and quantity of free network, scientific and, for us, development software, make it a logical choice for research work of this type. It seemed to interest the public and judges that we used Linux instead of more mainstream options. Only one person had the nerve to ask us about its competency as an OS and they were swiftly dealt with!

The program may seem simple to more advanced programmers, but our conclusion was that if two 17-year-olds and a 16-year-old can achieve simple computer learning on our first attempt, then more advanced programmers out there can surely do better. Some form of artificially created intelligence must surely be part of our future.

Future plans...

The project may not end here. We have been invited to display the project again at the Young Innovators Exhibition in Belfast later in the year. As we all plan to study some sort of computing at third level, this work may go on for some time yet. **LXF**
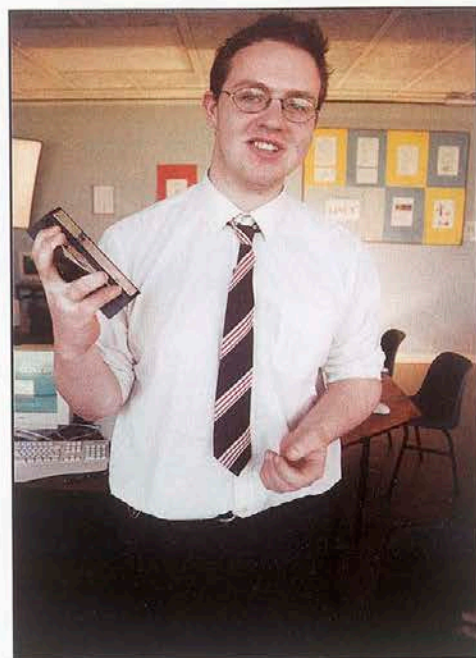
Careful with that award, Darius.

This is Alex, who seems to deserve the award the most, for having to put up with the other two.

LINUX USER GROUPS

Wherever you are in the world, there'll be a Linux User Group somewhere near you. There are thousands of dedicated User Groups all over the UK alone, so find the one nearest to you now!

UK LUGS

One of the great things about Linux is that you are never alone. There are thousands of User Groups, full of members keen to help with your problems, discuss ideas and generally natter about all things Linux.

We have collected a load of information here so you can find the LUG closest to you.

You can find lots more information online at:

www.linuxformat.co.uk/lugs or at **www.lug.org.co.uk**

1 HAMPSHIRE
URL: www.hants.lug.org.uk
Contact: Ken Adams

2 BRISTOL & BATH
URL: www.bristol.lug.org.uk
Contact: Dave D

3 SCOTTISH
URL: www.scottish.lug.org.uk
Contact: Tony Dyer

4 OXFORD
URL: www.oxford.lug.org.uk
Contact: Alasdair G Kergon

5 BROMCOM (KENT)
URL: www.brighton.lug.org.uk
Contact: John Mills

6 BRIGHTON
URL: www.hants.lug.org.uk
Contact: Johnathan Swan

7 SUSSEX
URL: www.sussex.lug.org.uk
Contact: Mike Pedley

8 NORTHANTS
URL: www.northhants.lug.org.uk
Contact: Kevin Taylor

9 ANGLIAN
URL: www.anglian.lug.org.uk
Contact: Martyn Drake

10 MILTON KEYNES
URL: www.mk.lug.org.uk
Contact: Denny De La Haye

11 DONCASTER
URL: www.doncaster.lug.org.uk
Contact: Andy Smith

12 SOUTH EAST
URL: www.southeast.lug.org.uk
Contact: Ian Reason

13 WEST WALES
URL: www.west-wales.lug.org.uk
Contact: Ken Adams

14 WOLVES
URL: www.wolves.lug.org.uk
Contact: Jono Bacon

15 PETERBOROUGH
URL: www.peterboro.lug.org.uk
Contact: Steve Gallagher

16 EDINBURGH
URL: www.edinburgh.lug.org.uk
Contact: Alistair Murray

17 TYNESIDE
URL: www.tyneside.lug.org.uk
Contact: Brian Ronald

18 LEICESTER
URL: www.leicester.lug.org.uk
Contact: Clive Jones

19 GREATER LONDON
URL: gllug.linux.co.uk
Contact:

20 SURREY
URL: www.surrey.lug.org.uk
Contact: James Chivers

21 CAMBRIDGE
URL: www.cam-lug.org
Contact: Unknown

22 DEVON & CORNWALL
URL: www.lug.termisoc.org
Contact: Alex Charrett

23 ESSEX
URL: www.epos.demon.co.uk
Contact: Unknown

24 MANCHESTER
URL: www.manlug.mcc.ac.uk
Contact: Ted Harding

25 WEST YORKSHIRE
URL: www.scs.leeds.ac.uk/wylug
Contact:

26 WEST YORKSHIRE
URL: www.wylug.lug.org.uk
Contact: Nigel Metheringham

27 SHEFFIELD
URL: www.sheflug.co.uk
Contact: Richard Ibbotson

28 STAFFORDSHIRE
URL: linux.ukweb.nu
Contact:

29 NORTH-EAST
URL: www.shofar.freeserve.co.uk/NELUG

30 LONDON
URL: www.lonix.org.uk
Contact:

31 THAMES VALLEY
URL: www.sclug.org.uk
Contact: Nick Lambert

32 LIVERPOOL OpenSource
URL: linux.liv.ac.uk/LIV_LINUX_UG
Contact: Simon Hood

LUG ORGANISERS!

If you're not listed here, or we have your details wrong, please contact us. It would help if you could fill in the form here, or on our website at www.linuxformat.co.uk/LUGs

Name of LUG: _____
Location: _____
Contact Name: _____
Website address: _____
Regularity of meetings: _____
Any other information: _____

Send the form to:
LUGS!, Linux Format,
30 Monmouth Street,
Bath, BA1 2BW,
or email your details to:
linuxformat@futurenet.co.uk

WORLDWIDE LUGS

AUSTRALIA:
ADELAIDE LUG
Url: www.linuxsa.org.au
Contact: MatthewTippet
Email: mtippet@anu.edu.au

PERTH
Url: plug.linux.org.au

AFRICA
STELLENBOSCH
Url: www.entropy.sun.ac.za/
Contact: Abraham vd Merwe
Email: ixion@entropy.sun.ac.za

EUROPE
GHENT
Url: llsgg.rug.ac.be/
Contact: Wim Vandeputte
Email: wvdputte@llsgg.rug.ac.be

GOTHENBURG
Url: nain.oso.chalmers.se/LUGG/index.html

LISBON
Url:
www.students.iscte.pt/ · a12593/gul.html
Contact: Paulo Trezentos
Email: Paulo.Trezentos@iscte.pt

AUVERGNE
Url: www.linux-arverne.org/
Contact: Cyril Hansen
Email: Cyril.Hansen@wanadoo.fr

SOUTH AMERICA
SAO PAULO
Url: gul.linux.ime.usp.br/
Contact: gul@ime.usp.br

BUENOS AIRES
Url: N/A
Contact: Daniel E. Coletti
Email: dcoletti@impost.com.ar

MONTEVIDEO
Url: www.linux.org.uy/
Contact: Rodolfo Pilas
Email: uylug@linux.org.uy

LIMA
Url: linux.unired.net.pe/
Email: linux@unired.net.pe

NORTH AMERICA
ALASKA
Url: www.aklug.org/index.html
Contact: Dee McKinney
Email: deem@wdm.com

LOS ANGELES
Url: www.lalugs.org/
Contact: Dan Kegel
Email: dank@alumni.caltech.edu

BAY AREA
Url: www.balug.org/
Contact: Arthur F. Tyde III
Email: aftyde@balug.org

DENVER
Url: spot.elfwerks.com/ · clue/
Contact: Lynn Danielson
Email: lynnd@ihs.com

TAMPA
Url: terrym.com/slug/index.html
Contact: Paul Foster
Email: paulf@quillandmouse.com

BATON ROUGE
Url: www.brlug.net/
Contact: Dustin Puryear
Email: dpuryear@usa.net

VIRGINIA Tech
Url: corvette.me.vt.edu/pages/index.html
Contact: Bucky LaDieu
Email: nega@vt.edu
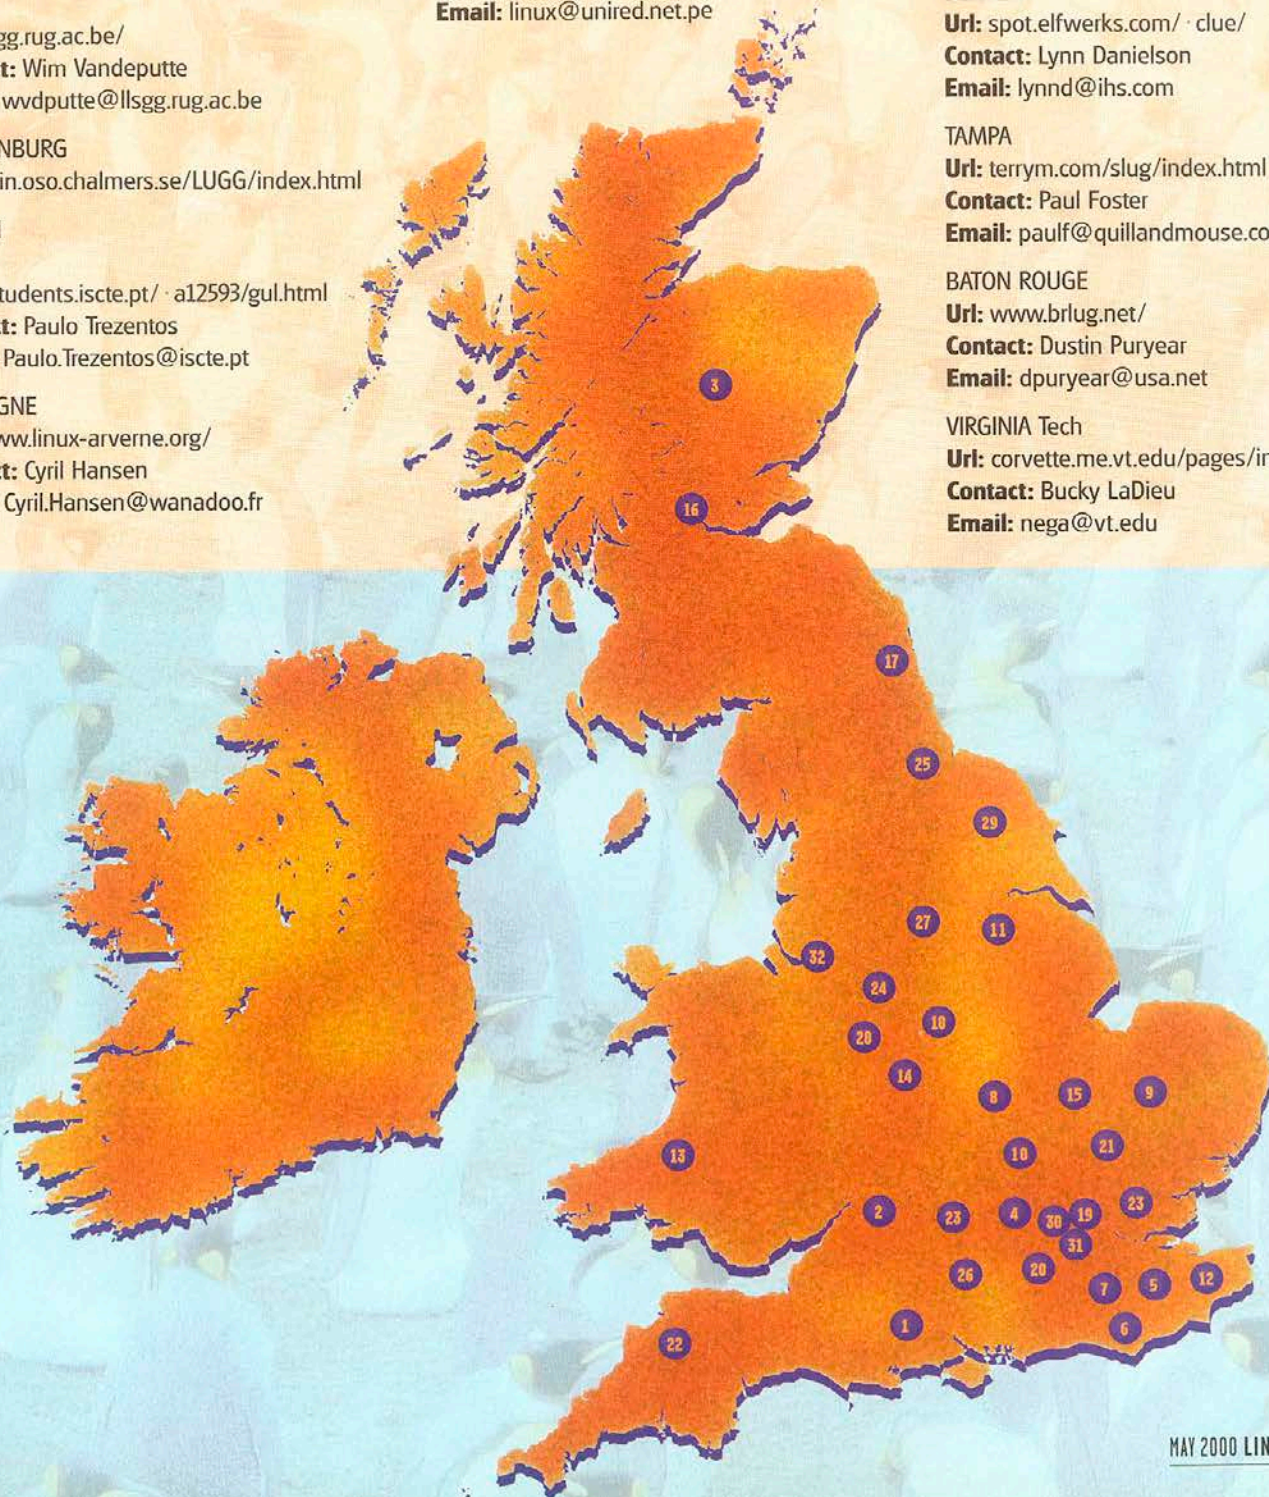
Licensed to THRILL!

Paul Talacko looks at the hot issue of licensing and what you can and can't do with your Linux software.

A dry topic like copyright law rarely arouses great passion, but if you want to start a flame war between Linux users, you could do worse than mention copyright law and the issue of licensing. Heated debate arises because of some fundamental differences between the ideas of the 'free software community', the 'open-source community' and how traditional software businesses operate. It goes to the very heart of what you can and can't do with a program.

Copyright law is often regarded as the bane of the free software community. It allows only the 'author' of software to make copies unless agreed otherwise. Businesses use this law to make money. For example, when you buy a product from a company like Microsoft, you are agreeing to a licence that does not allow you to copy the program, and you most certainly don't get to see the source code.

Open-source software and, in particular, the idea of copyleft, takes this concept and turns it on its head; in effect, using the law of copyright against itself. Below we'll have a look at a few of the common open-source licences and see what ignites the flame, but first we have to know what 'free' and 'open-source' really mean.

Free or open-source

The concept of freedom is essential to the idea of how Linux is licensed, but apply the same word to software and there's no end of confusion. Free, in the sense we are talking about, doesn't mean that you don't pay money for the software; rather you can do with it what you will. It is unfortunate that we all speak English because the language doesn't have an unambiguous word to describe the term. In Japanese it is 'jiyuu'. In French it is 'libre'. In a recent EU discussion paper on free software (**http://eu.conecta.it/**), the Franglais 'libre software' was used to make the idea clear.

This sort of freedom is a matter of principle, says Richard Stallman, founder of the GNU project and the Free Software Foundation. Mr Stallman's attitude on freedom and on the licensing of software has caused quite a bit of bad blood between the free software and open-source movements.

"I am disappointed with the values that the Open Source Movement appeals to and promotes, because they are limited to practical issues – how to get powerful and reliable software," he said. "In effect, they assign freedom and community no value in themselves. For me, freedom and community are even more important than powerful and reliable software. The main difference

between the Open Source Movement and the Free Software Movement is in the philosophy; by comparison, the difference in which licences we accept is small."

The freedoms that Mr Stallman refers to are as follows:

1. The freedom to run the program for any purpose.

2. The freedom to study how the program works and the ability to adapt it to your needs.

3. The freedom to redistribute copies so you can help your neighbour.

4. The freedom to improve the program and release your improvements to the public so that the whole community can benefit.

The open-source community, on the other hand, accept licences that comply with the Open Source Definition (**http://www.opensource.org/osd.html**).

It is very similar to the four freedoms and the differences are subtle. Under the OSD you must allow free redistribution of the program, the source must be made available, and modifications must be allowed to be made and distributed (although there is a limit in that these may be distributed as patches). Furthermore, the licence must not discriminate against persons or fields of endeavour, and must not contaminate any software it is distributed with.

These ideas are distilled in the licences, and in particular the GNU General Public Licence, the archetypal open-source licence. It also introduces something called 'copyleft'.

The GNU General Public Licence, the GNU Lesser Public Licence and Copyleft

When he was working in academia, Mr Stallman lived in an atmosphere of mutual assistance. Source code was shared freely and help was afforded whenever it could be. Then most of these bright students got snapped up by software companies and could no longer share the fruits of their work. Mr Stallman single-handedly decided that software should be shared and programmers should help one another. He started the Free Software Foundation and the GNU project, with the aim of having all the software needed to run a computer for free. But that wasn't the end of the matter. Mr Stallman knew that it would be legal for companies to take his work and make it proprietary.

"I knew that this would be legally possible if I did not do something to block it. And I figured that if it were legal, it would happen," he said.

This is the essence of what makes the GPL so important, but also different from other open-source licences. Not only do modifications to the program have to be subject to the GPL, but any program that links in and uses it must be as well. This is called copyleft and it has caused a few problems, especially with libraries.

A library is a group of code that may be used by several programs. The GPL would prohibit non-free programs from using a library. For example, *WordPerfect* is a non-free program, but it uses the GNU C libraries to run. It wouldn't be allowed to work on your Linux system if it was subject to the GPL. The Free Software Foundation therefore has another type of licence, the LGPL. This used to stand for the Library General Public Licence, but is now the Lesser General Public Licence. Richard Stallman writes that the LGPL should only be used where the features →
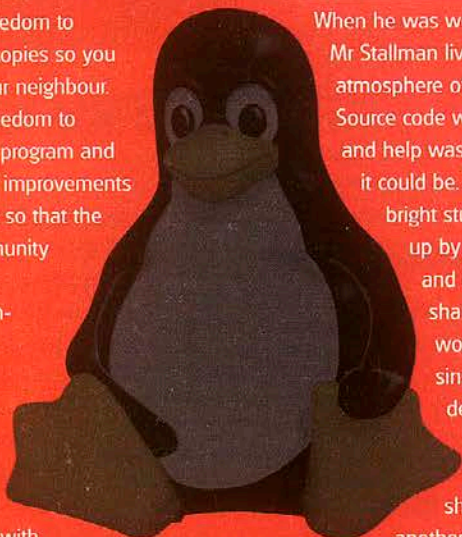
➔ of the libraries are available in non-free libraries. Where the features of a free library outstrip those of a non-free library, making it subject to the full GPL rather than the LGPL will force many proprietary software manufacturers to make their programs free.

Because free software like Linux source code is open, there is ample scope for people to make different, incompatible versions of it, but the GPL forces all amendments to be freely available. In effect, this means that if the amendments are any good, they will be incorporated into the main body of the program. And if they're bad? Well, no one will use them anyway.

Mozilla and Netscape Public Licences

The world was shocked when Netscape Corporation gave away the jewel in their crown, *Netscape Communicator*, the web browser, email client, news client and web page editor.

The source code for *Communicator* was handed over to the Mozilla Project (**http://www.mozilla.org**). To protect their interests, Netscape didn't release the code under the GPL, but under two licences: the MozillaPublic Licence (MPL) and the Netscape Public Licence (NPL).

With the MPL, Netscape had to come up with a licence that would satisfy all-comers. It had to suit the open-source developers who would actually be writing the code in their spare time, and it had to be compatible with the desires of

corporations – in particular Netscape themselves, who would have to be able to use the code in some way that would allow them to make a profit. The MPL allows you to: "use, reproduce, modify, display, perform, sublicense and distribute the Original code (or portions thereof) with or without Modifications, and/or as part of a Larger Work". However, when it is incorporated into a larger work, that larger work does not have to be released to the public. This is important for Netscape, because it means that the Mozilla code can be taken and turned into *Communicator* with proprietary third-party software being added to it.

Both of these licences are different from the GPL in two subtle respects. Firstly, it is not a real copyleft because amendments which are not made to the original code but which are in a separate file don't have to be publicly released. Secondly, it is not compatible with the GPL. However, you can incorporate GPL code with a small modification to the GPL.

The terms of the NPL and MPL are almost the same, except that the NPL allows Netscape Communications to include something called Netscape Branded Code that does not have to be released back to the community. It is therefore not a 'free' software licence.

Artistic Licence

The Artistic Licence allows the authors of certain Perl packages to retain some control over their work. People can modify a program, provided that such modifications are clearly distinguishable from the original version. Modifications must be made freely available, unless they are only used within an organisation. The program may also be included in commercial distributions, provided it is not 'visible' to the outside world.

> **There is, therefore, an argument that the QPL is freer than the GPL.**

KDE and the Qt Public Licence

If you use Linux, you probably use the *K Desktop Environment. KDE* is based on a series of libraries called *Qt* that were written by Norwegian Company Troll Tech. The latter wanted to keep some control over *Qt* and came up with the Qt Public Licence. This allows distribution of modifications to *Qt*, provided the changes are distributed as patches (ie, they're not part of the main body of the code), and it's up to Troll Tech to incorporate these modifications back into the body of *Qt*.

This is markedly different from the GPL, which allows all modifications. It also means that the QPL is not free, in the GNU sense. As a result, many people prefer to use the *GNOME* Desktop (http://www.gnome.org) as a matter of principle, simply because it is subject to the GPL. Troll Tech also allow companies to make proprietary software where the code is not released at all. There is, therefore, an argument that the QPL is freer than the GPL.

Matthias Kalle Dalheimer, one of *KDE's* developers, says,"My company does custom *Qt* development, and I do *KDE* development in my spare time. If *Qt's* combination of an open-source licence and a commercial licence didn't exist, I would have to use different toolkits for my free and my commercial

Bill Joy, the father of Java, which is distributed under the Sun Community Licence.

Matthias Kalle Dalheimer, one of the developers of KDE.

programming — something that is virtually impossible to do. Thus, *Qt*'s licences give me the freedom to choose how I want to license my own programs, something that the GPL doesn't."

The BSD and X Licence

These are very simple licences that apply to a Linux-like product called BSD and the X Windows System. They allow others to modify and use the source code, but don't require them to release it back into the community. This is why Apple chose BSD as the basis for their new operating system, Mac OS X: they can incorporate the code for their crown jewel, their user interface, without allowing others access

to it. The Free Software Foundation recommends that developers don't use this licence.

The Apple Public Source Licence

This licence is one under which Apple have released parts of OS X and some of *QuickTime*. It is free-ish, but not, according to Richard Stallman, nearly free enough. Unlike the GPL, there are some restrictions on the source code. In particular, the licence does not allow you to make modifications for your own use, and therefore disrespects privacy, requires all modifications to be notified to Apple and allows Apple to revoke the licence any time someone makes an accusation of patent or copyright infringement.

Sun Community Licence

Then we come to Sun. Sun make some wonderful products, like the Solaris operating system which still out-performs every other (even Linux). This can be licensed under the Free Solaris Software Licence and allows researchers and customers to view the Solaris source code, but they may not share it. This is not a free licence.

Other products from Sun, like Java and *StarOffice*, the office suite, are being released under another licence, the Sun Community Source Licence. This is not a free licence, nor does it comply with the open-source definition. According to Danese Cooper of Sun Source Licensing, this licence is used to create a community of developers on

technologies which require careful stewardship. Under this licence, developers can keep their modifications private and can share code. However, they can only release their own products under certain conditions.

"Before productive distribution of any derivative works, licensees must sign a Commercial Use agreement with Sun and must pass the appropriate compatibility test suite," said Ms Cooper, so this is not a free licence.

Since February, Sun have a third type of licence for which they have requested open-source certification from the Open Source Initiative. This is the Sun Industry Standards Licence. "The licence requires that anyone who wishes to use Sun code in a way that deviates from the standard must publicly supply their specification and reference implementation," Ms Cooper said.

Finally, Sun are releasing one of their products under a free software licence. *Forte* for Java, Community Edition, an IDE (Integrated Development Environment, a program which makes writing software easier), is being released under the Mozilla Public Licence.

These are the major licences. There are others, like the JavaScript licence, the Perl licence and the LaTeX licence. Life would be easier if there were fewer, but as long as copyright law is not changed and companies want to release source code while retaining some control over their software, there won't be fewer.

As we've seen, some licences, like the BSD licence, allow free code to be incorporated in non-free products, while the NPL allows non-free code to be incorporated into *Communicator*. The QPL, on the other hand, allows Troll Tech to keep control over their code.

However, Richard Stallman's stance on what is acceptable means that the GNU General Public Licence is the standard open-source licence. Not only does it make it illegal to ever make the code secret, but it introduces the idea of copyleft, which makes sure that source code will be free for ever. **LXF**

The GNU General Public Licence is the standard open-source licence.

Richard Stallman, a leading proponent of software 'freedom'.

GPL Do's and Don'ts

Most of the software on your Linux box will be subject to the GPL, so here is a list of the basic do's and don'ts:

◆ **DO provide copyright notice if you copy and distribute source code.**

◆ **DO ensure modified source code has notices that you modified it.**

◆ **DO allow people access to the source code if you are only distributing executable versions.**

◆ **DO get permission of the copyright owner (which will often be the Free Software Foundation) if you want to change the terms of a modification.**

◆ **DON'T try to modify the terms of a GP program.**

◆ **DON'T try to restrict the use of any GP program.**

◆ **DON'T try to use parts of a GP program in a program that is not GPL.**

A FEW INTERESTING WEBSITES

For the GNU software philosophy:
http://www.gnu.org/philosophy

The text of the GPL:
http://www.gnu.org/copyleft/gpl.html

The Open-Source Definition:
http://www.opensource.org/osd.html

The text of the Artistic Licence:
http://www.perl.com/pub/language/misc/Artistic.html

The text of the Mozilla Public Licence:
http://www.mozilla.org/MPL/MPL-1.1.html

The text of the Netscape Public Licence:
http://www.mozilla.org/MPL/NPL-1.0.html

The Debian Free Software Guidelines and Social Contract:
http://www.debian.org/social_contract.html

The Apple Public Source Licence:
http://www.publicsource.apple.com/apsl/

Sun Community Source Licence:
http://www.sun.com/981208/scsl/principles.html

The X11 licence:
http://www.xfree86.org/3.3.3/COPYRIGHT1.html

What on earth is TeX?

In the first of our regular guides to useful Linux software, **John Palmer** explains the myth and mystery surrounding the thing they call TeX – he'll even tell you how to pronounce it.

When you unwrap your Linux CD-ROM from Red Hat, or your chosen distributor, you may not realise what a bundle of goodies is included. Indeed, one of the great packages that may not be immediately obvious is the typesetting system *TeX*.

IN A SENTENCE, WHAT ON EARTH IS IT?

The function of *TeX* is to take a text – a sequence of words, probably interspersed with typesetters' instructions like 'paragraph' and 'section-heading' and 'bold-face' – and turn it into an elegant, printed document.

HOW ARE YOU SUPPOSED TO PRONOUNCE THAT?

TeX is pronounced 'tech', as in 'technical', by decree of its originator, Don Knuth. Better to avoid arguments and get used to this, and never rhyme it with 'sex', even though there is a book called *The Joy of TeX*. Where possible it is written as a distinctive logo, with the 'e' dropped a bit between the T and the X; where this isn't so easy to do, for instance in email, you can simply write 'TeX'.

WHERE DID IT COME FROM?

TeX is particulary suited to typesetting mathematical documents.

TeX goes back to the late '70s when Donald Knuth, a professor of computer science at Stanford University, California, was working on his great opus, *The Art of Computer Programming*. He soon discovered that commercial typesetting made his work look ugly and sometimes even misrepresented his original ideas, so he set about making a program that would produce masterpieces of the typographer's art.

TeX had to be superb at handling mathematics. Indeed, maths is probably the hardest challenge a typographer has to face. In the days when type was set by hand, mathematical copy was charged at premium rates due to the complicated positioning, special characters and the ease with which a small slip can lead to a complete change in the meaning of a mathematical expression – a change very noticeable to a mathematician but not obvious to the ordinary reader, including perhaps the compositor himself!

TeX also had to be capable of setting very large documents, typically complete books of several hundred pages. It had to make sensible decisions about hyphenating words, justifying paragraphs, starting new pages and finding places for figures and tables in the text. It also had to make tables of contents, lists of figures and indexes and prepare them elegantly and in an adaptable fashion.

SO TEX IS A WORD PROCESSOR?

Not really. A word processor is essentially a very sophisticated typewriter; what you key in appears on your screen in as near as possible its final printed form – a WYSIWYG display (What You See Is What You Get). *TeX* is more ambitious than this as it actually attempts to do the job of a highly skilled compositor and layout-artist. If you want high-quality typesetting, you can't expect one program to do the text-editing, the screen-display and the output to paper. Updating the screen as you type is a difficult business and not really compatible with the best quality typesetting, so *TeX* doesn't attempt WYSIWYG editing.

SO TEX IS A DTP PROGRAM?

Again, not really. DTP programs are good at making up small documents where layout is all, but they aren't so hot on paginating long documents; you usually have to decide what to put on what page yourself. *TeX* actually excels on long documents. Every time you run *TeX*, it reviews all its decisions about where to break pages, where to put floating objects like figures and tables and where to break lines and hyphenate words. This review takes in presentation of the entire finished document, not just one page at a time.
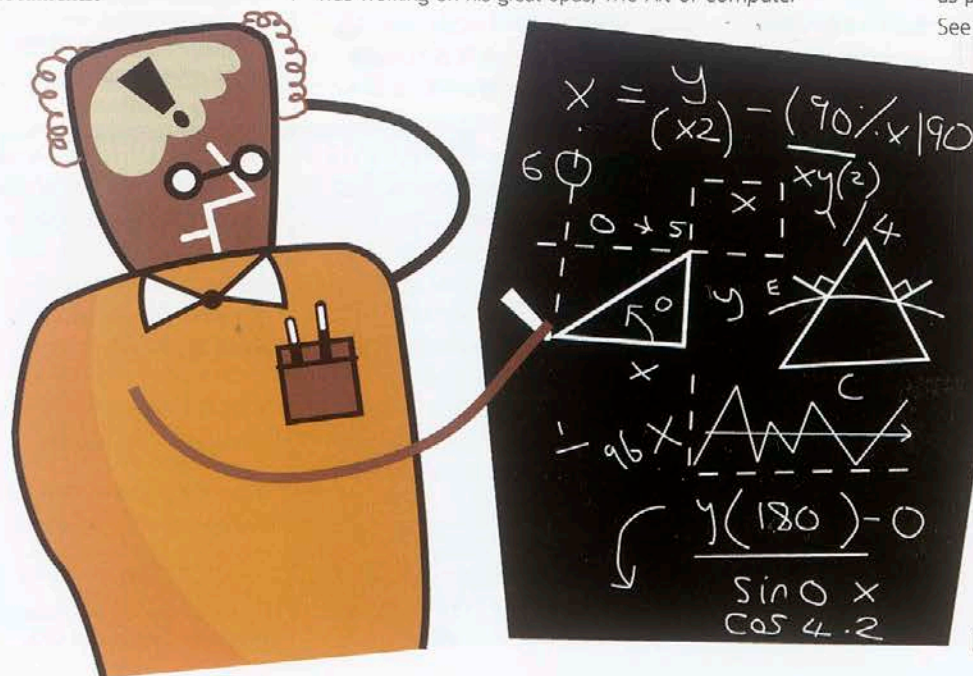
IS IT REALLY WORTH CHANGING TO TEX?

Well, if you're used to *Word* or *StarOffice*, you're happy with the typographical style of these programs and you're writing nothing larger than a long letter, or more complicated than an advertising flyer, there's probably no point in changing to *TeX* because it'll take you longer to learn *TeX* commands than it will to write your whole document!

If, however, you want to write a long document which includes tables of contents and indexes, or you want to use lots of mathematical notation, or you simply like the fonts that *TeX* makes available to you, it will definitely be in your interest to give it a go and see how good it is at producing elegant typeset documents.

THE TEX WAY OF COMPOSING A DOCUMENT

With *TeX* you don't edit a screen version of the printed page. Instead, you edit an input-file which contains all the words you're going to print, and all the directions about how you want them printed — paragraphs, section-headings and so on— but looks nothing like the printed page.

After you've made some changes to this input-file, you tell *TeX* to process it and then you view on the screen how it is going to look when printed. Every time you run *TeX* it processes the whole input-file, and so it can take radical decisions about where to break between pages, where to place floating figures and tables and so on.

With *TeX* you can use any text editor you like. On Linux, you'll probably use *vim* (*vi*), or maybe *emacs* with a *TeX*-specific format such as *AucTeX*. You can even use your favourite word-processor if you remember to make it save files as ASCII text. However, word-processors are a bit clumsy for this job, simply because they are more than just text-editors.

WHAT IS THIS LATEX THING?

TeX itself has about 300 basic typesetting operations, but is frankly impossible to use without a collection of higher-level commands called macros, which constitute the package called *Plain TeX*. *LaTeX* is a further collection of macros at an even higher level, nearer to the level of abstraction at which authors commonly think of their text.

The 'LA' part of its name comes from their original author, Leslie Lamport. More recently it has been enhanced by an enthusiastic group of volunteers who run the *LaTeX3* project. The current version of *LaTeX* is called *LaTeX2e*.

LATEX IS AIMED AT AUTHORS

When you use *LaTeX* you don't generally think much about typeset layout; instead you concentrate on writing what you mean, but interject a few commands like:

> **"\section", "\emph" and "\begin{table}"**

As you write, you avoid thinking too much about presentation of the words. You can control that by a declaration like:

> **\documentclass[a4paper,11pt]{article}**

at the beginning of your text. Changing the style or typeface of an entire document is just a matter of changing a couple of lines at the start of the input-file.

As most users of *TeX* regularly use *LaTeX* in some form, you should understand that from here on that what I say about *TeX* also applies to *LaTeX*.
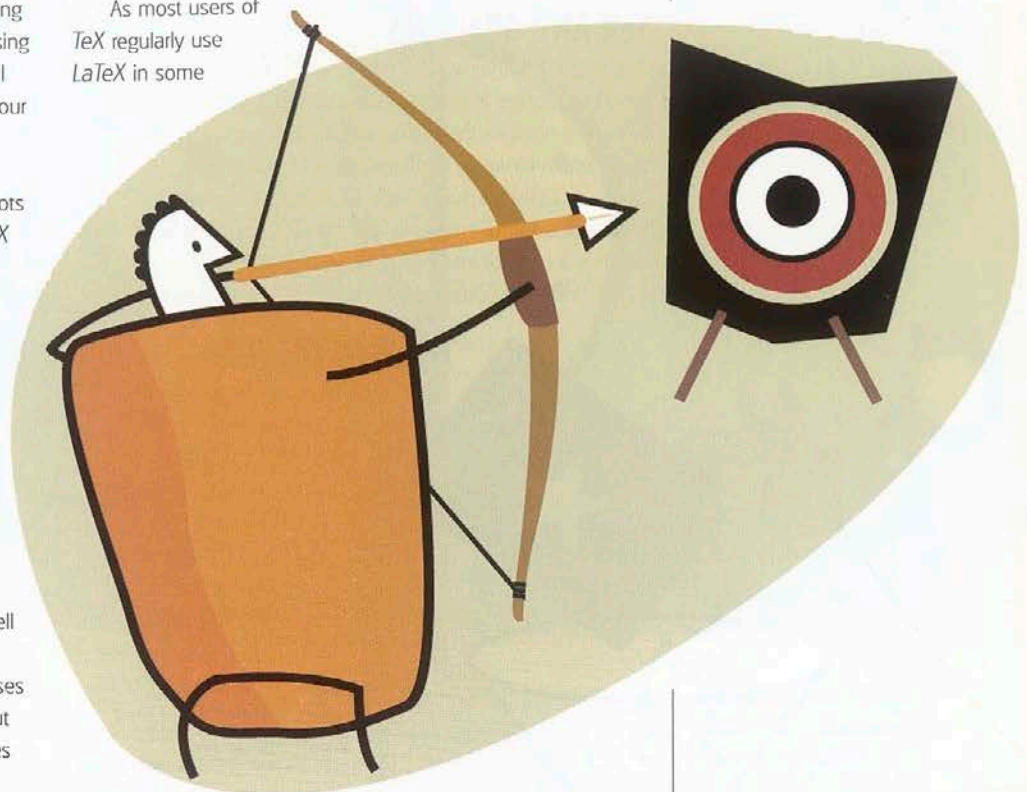
LaTeX is aimed at authors and can be used much more naturally than TeX alone.

TeX AND PRINTERS (machines, that is)

When *TeX* processes an input file it generates three outputs. The output on the screen tells you roughly what's going on and asks for human help if *TeX* encounters a problem it can't solve by itself. The log-file contains most of what appears on the screen, plus additional detail which helps experienced users of *TeX* to make it perform a little better. The device-independent file (with a name suffixed ".dvi") contains the instructions for making the printed output, either on the screen, or on any kind of printer.

To make the actual output, you have to process the .dvi file with a driver program specific to the output-device. For this, you usually need at least two drivers — one for your screen and one for your printer. The screen driver most used on Linux is xdvi, which displays a *TeX*-set document on X Windows. The printer-driver may be specific to a particular make of printer, such as dvilj which drives a Hewlett-Packard printer using its native control language PCL, or it may be something that's more generally applicable.

TeX AND POSTSCRIPT

One of the most used drivers (particularly on Linux) is dvips, which translates the .dvi file to PostScript, Adobe's page-description language. This can be sent direct to those printers able to interpret PostScript directly, or interpreted by *Ghostscript* for simpler printers that can't. There are great advantages in the PostScript method of printing as you can embed images into →

→ the page, especially if you use *PStricks*. There is also a variant of *TeX* called pdfTeX, which bypasses the .dvi stage and generates Adobe Portable Document Format (PDF) directly.

TeX AND GRAPHICS

TeX originally came without any serious support for including pictures in your typeset pages, but Knuth foresaw a need for this facility and the process has been made much easier by the introduction of PostScript. There are *LaTeX* packages such as *PStricks* which allow for the inclusion of images in *TeX*'s output, and a book which is a mine of information on the subject, *The LaTeX Graphics Companion*.

TEX AND THE WEB

This is probably the fastest growth area in *LaTeX* facilities. There are tools for converting between *LaTeX*, plain *TeX* and HTML, XML, etc, and the *LaTeX3* project seems to be moving towards the option of SGML-like notation as input to *LaTeX*. The best guide to this is a 1999 book which is called *The LaTeX Web Companion*.

TEX AND METAFONT

TeX had to come with some elegant typefaces. Don Knuth didn't like the offerings from usual type-foundries, so he designed his own, including the elegant Computer Modern (CM) type-family. Being a mathematician, he designed it mathematically, and made his own system for creating new characters and typefaces by computer. This is *METAFONT*.

Most users of *TeX* and *LaTeX* don't have to learn how to use *METAFONT*, but it is part of every decent *TeX* distribution, and when you use *TeX* it does some important work behind the scenes. If, however, you're interested in constructing new fonts, typefaces for exotic languages or strange new symbols, *METAFONT* is your tool.

METAFONT lets you design each new character by specifying the co-ordinates of the salient points of the letter and a few curves to join them up. What makes it exciting is that you base each letter on a common set of parameters that define the characteristics of the typeface. Then, by varying the values of

The PStricks package makes including images in TeX's output that much easier.

these parameters, you can change the style of the typefaces without redesigning the individual letters. Hence it only involves a few changes to take an existing Roman face and create a companion bold face or italic, or even a new family type entirely.

TEX AND NON-LATIN SCRIPTS

The versatility of *METAFONT* has encouraged many people to design fonts for various non-Latin alphabets and other scripts; not only the obvious Greek and Cyrillic, but also non-European alphabets like Thai, Devanagari and several other Indian scripts.

TeX can handle languages which are written from right-to-left, so it can also set Arabic and Hebrew, and even the ideographic scripts of Chinese, Japanese and Korean. Many of these fonts are available in *TeX Live* and the CTAN Archive, which are described later.

TEX AND GNU

From the start, Don Knuth distributed *TeX* free of charge, in much the same spirit as Richard Stallman and the other contributors to the GNU project. In fact, GNU, the project to create a free Unix-like system, came after *TeX* and adopted *TeX* as its regular typesetting system in place of Unix's rather dated *troff*.

Much *TeX* and *LaTeX* software is now released under the LaTeX Project Public Licence (LPPL), which is very similar to the GNU Public License (GPL).

THE TEX ARCHIVE

All the TeX software you are ever likely to need is available for free download from CTAN, the Comprehensive TeX Archive Network. This consists of three principal servers – in Germany, the UK (Cambridge University) and the US – and a host of satellites that mirror these.

CTAN is maintained by a small and selfless group of volunteers who ensure that the three principal servers all contain the same data. For users in Britain, the easiest way to inspect
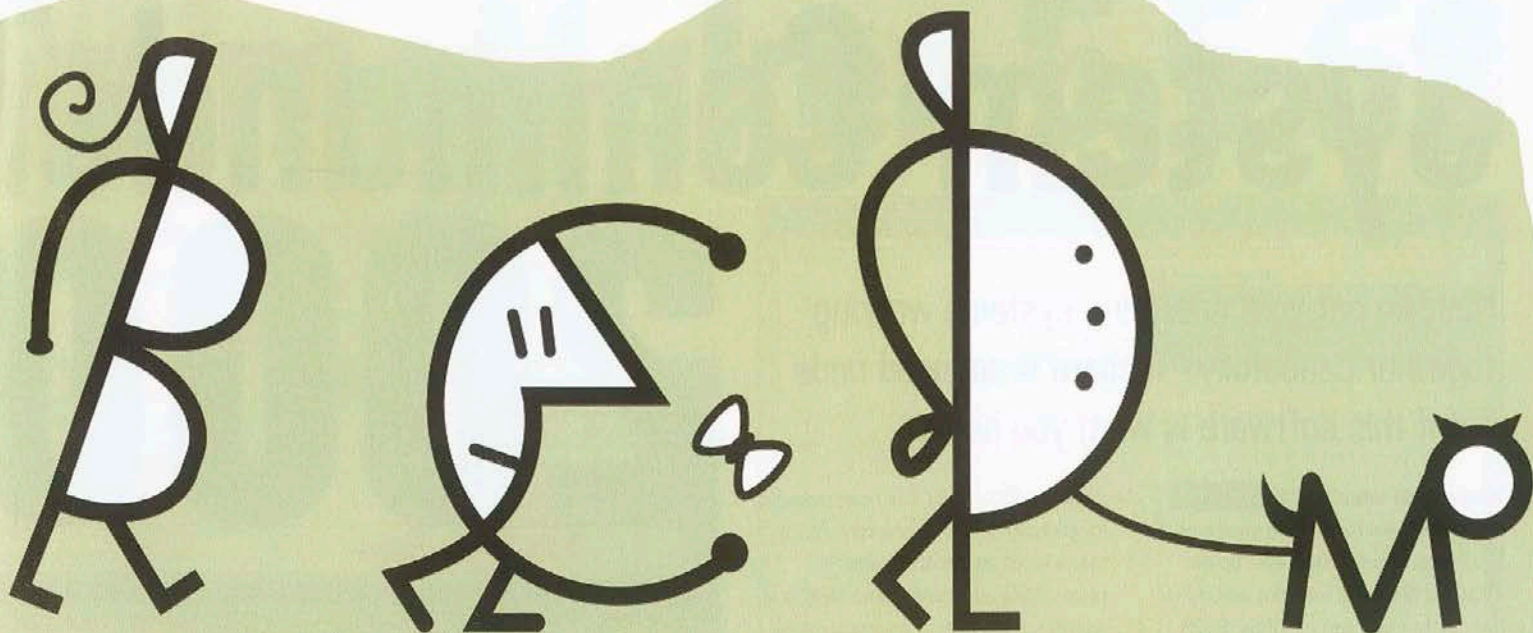
the contents is to point your web browser to the following website: **http://www.tex.ac.uk/**

TEX USER GROUPS

It's a very good idea to join one of these as you may get not only a journal, but also access to discounted books and free or inexpensive software distributions like *TeX Live*, which is not unlike the *TeX* of Red Hat Linux but is much richer in special formats, new fonts, and other goodies.

The principal user group is TUG (TeX Users' Group) which is US-based and publishes the journal *TUGboat*.

There are also many other national groups. The British group UKTUG has a journal, *Baskerville*, and also maintains a useful document called The TeX FAQ, which answers many of the questions *TeX* novices (and many experienced TeXnicians) ask.

The key source of information is **http://www.tug.org/**

A great place to get help on using TeX is the Usenet discussion group, comp.text.tex.

which will refer you to all the other sources; UKTUG has its own page at **http://uk.tug.org/** and you can also email enquiries to **uktug-enquiries@tex.ac.uk**

TEX DISCUSSION GROUPS

There are several, but the new user is going to be most interested in the liveliest source of advice on using *TeX*, and this is undoubtedly the Usenet newsgroup comp.text.tex.

For those people who do not have easy access to Usenet news, there may soon be a way of getting access to this *TeX* group by using email.

WHERE TO READ MORE ABOUT TEX

The essential guide, required by all serious users of *TeX*, is Don Knuth's *The TeXbook*. You are also going to need a book on *LaTeX*, of which there are lots. Leslie Lamport's *LaTex, a document preparation system*, is particularly authoritative because it was written by the author of *LaTeX*, itself. Make sure you get the second edition (1994) which is up-to-date with the improvements introduced in *LaTeX2e*. After that you may need *The LaTeX Companion*, *The LaTeX Graphics Companion* and even the latest in the series, *The LaTeX Web Companion*. All these books are (naturally) typeset with *TeX*, they're all published by Addison-Wesley, and UKTUG (among other *TeX* user-groups) can get them for their members at discounted prices.

John Palmer is a member of the committee of UKTUG.

TeX comes with its own verstile font description system, METAFONT.
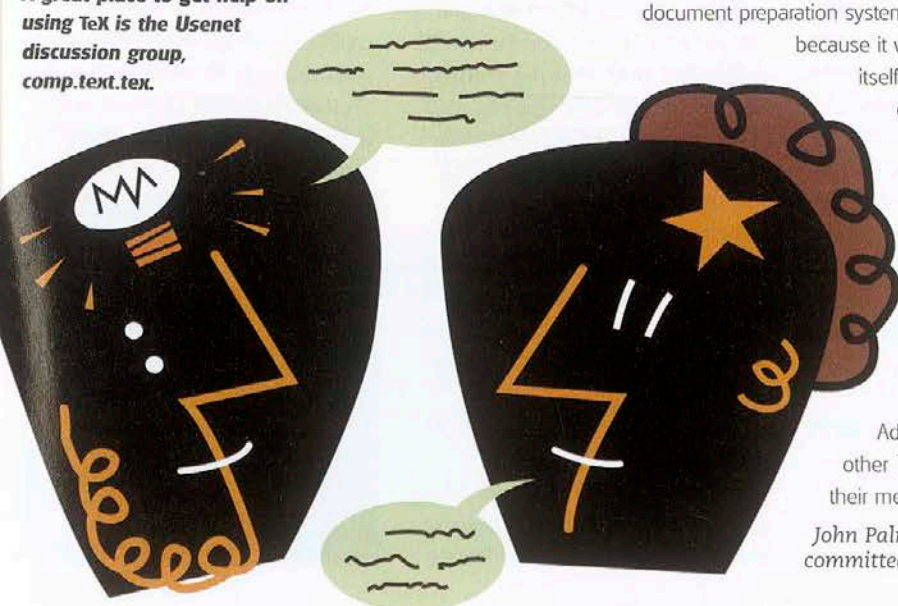
System Commander 2000

Need to get your operating systems working together peacefully? **Richard Drummond** finds out if this software is what you need...

Contrary to what the US DoJ might claim, we now have more choice of PC operating systems than ever before. Okay, so Windows may not entirely fit the bill, but you can try Linux, BeOS, OS/2 and uncountable flavours of Unix. The obstacle that stops many users from taking the alternative OS plunge is the difficulty in getting competing operating systems to co-exist on one machine. *System Commander 2000* helps to overcome this. It is an OS manager and provides a smart OS installation wizard, a GUI-based partitioning tool and a boot menu.

A la carte

SC2000 is heavily biased towards Windows and DOS, but it does support Linux. Installation must be performed from either DOS or Windows, though. The boot menu gets installed into the boot record of your first drive and the rest of the *SC2000* software (the OS Wizard and the partitioning tool) lives in the first primary FAT partition. This

takes up about 3Mb. The boot menu is functionally similar to any other boot menu, such as the NT *Loader* or Linux's *LILO*, albeit with more bells and whistles. It provides a graphical menu in a faux Windows95 style, with each entry optionally illustrated with an icon.

Selection may be made via the keyboard or the mouse. You get the standard level of configurability, such as selecting a timeout period after which the default OS will load, and you can select which is the default or leave the default as the last OS booted.

Embellishments include a global password system, a screensaver and alert sounds. The OS Wizard, the partitioner and the preferences editor may all be launched from the boot menu. The menu knows how to start up all versions of DOS or Windows, but to boot into Linux you need *LILO* on your system, preferably installed into the superblock of the first partition. *LILO* doesn't need to know about other OSes you have on your system, though.

Making room

Installing a new operating system is not a task for the faint-hearted, but *SC2000*'s OS Wizard simplifies the procedure. You select which operating system you want to install and the wizard will analyse your system and give you a recommendation on how the OS should be installed. Typically, this will require creating new partitions and resizing old ones, and if so, *SC2000* will tell you. You can accept its suggestion or tweak the settings yourself. The Wizard will then perform any resizing operations and create and format the partitions necessary for
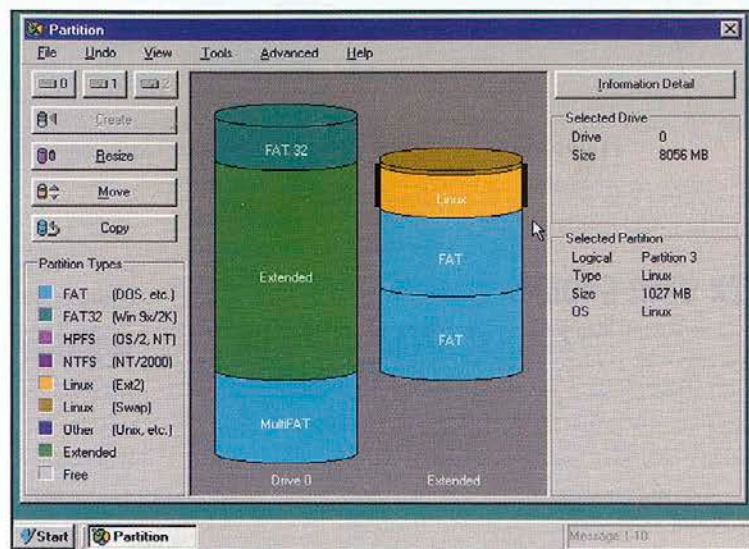
If you're scared by fiddling with your partitions, this software will help.

installing the new OS. Next time you boot your machine the boot menu will notice the changes to your drive and ask you if you want to add the new OS to its menu. When choosing to install Windows, DOS or OS/2, you may elect to install them in isolation or to share an existing FAT partition.

This is possible via *SC2000*'s so-called MultiFAT system. Each flavour of DOS, Windows or whatever gets installed in a unique directory in the root of your primary FAT partition, along with the necessary system files. When you want to boot one of these, the boot menu copies the COMMAND.COM, AUTOEXEC.BAT, etc, into the root directory and uses them to boot. When you next choose to boot a different OS from that partition, the system files are

backed up to that directory and the files for the new OS copied into the root and booted. The graphical partitioning tool allows you to manually modify the drive structure. You can create and delete any partitions, and non-destructively resize, move and copy FAT, FAT32, NTFS and Linux ext2 partitions. Each operation you perform can be reversed later. This applies to partitioning done by the OS Wizard too. **LXF**

FREE ALTERNATIVES

Until recently, all Linux distros were shipped with *FIPS*, a DOS tool to resize FAT partitions. This was enough to put many people off installing Linux for a start. Recent Red Hat releases have included the *Disk Druid*, a graphical replacement for *FDISK*, while Mandrake Linux has the experimental *Disk Drake* (see http://www.linuxmandrake.com/diskdrake) which can also resize partitions. Alas, none of these can measure up to *SC2000* in terms of reliability or features.

FrameMaker 5.5.6B

Publisher: Adobe **Web:** www.adobe.com

Price: Trial beta version

Adobe have been at the forefront of DTP since they invented *Postscript*, so it's no surprise that their publishing software is the first to make the jump to the Linux platform. *Acrobat Distiller* and the *Acrobat Reader* have been available for some time, but *FrameMaker* is the next logical progression.

FrameMaker is a WYSIWYG document layout package, but it isn't really aimed at the high-end publishing market. Its strength is in dealing with complex typographical requirements and heavily styled and structured text. Mmmm, you might think, this sounds a little like *TeX*.

Well, the people who might want to use it are certainly currently setting all their documents using *TeX*: there aren't very many choices for setting complex mathematical formulae, after all. I have to say that people who intuitively know *TeX* are probably better off sticking with it. It is powerful, and if all the arcane tags and macros are second nature, it would be a bit of a shock to have to construct formulae with point and click actions, or style text by selecting it with a cursor.

For everyone who isn't that intimate with *TeX*, *FrameMaker* opens the way to create some complex,

Basic DTP layout tools are provided – a simple newsletter would be easy to create in FrameMaker.

colourful and professional-looking documents. You can set complicated formulae using the built-in symbols, although you will have to flit around the various sections of the equation composer, selecting from symbols, operators, calculus functions and so on. This can take some time to get the hang of, but the amount of control over the output is good.

Like *TeX*, you can set up conditional text, place inline graphics and, stylewise, do just about do anything with text. *FrameMaker* isn't just restricted to print publishing either, since it has the ability to export HTML and XML documents too. This is done by linking a document's text styles to their HTML/XML counterparts, which is a pretty easy way to make sure that everything looks as you intended.

Other elements of the design will also translate to the web. Graphics output relies on third-party libraries and conversion utilities, which are included in the download. These should work fine, but you may experience difficulties if you already have different versions of these files installed.

BETA TEST

As we all probably know, there is a great wealth of software on general release which is actually still in a beta version. We want to be able to discuss this software and tell you about it, but it wouldn't be proper to review anything other than the official release version. Hence our Beta Test banner, which you'll find on articles concerning pre-release software. We will also not be scoring such software, but we will provide a verdict, indicating whether it is looking promising or not.

The equation editing system is not for the impatient, but it does give you very precise control over every element of complex formulae

$$\frac{1}{2}u^2 + \frac{P}{\rho} + gh$$

$$\frac{\delta u}{\delta t} + (u \bullet \nabla)u = -\frac{1}{\rho}\nabla p + \nu\nabla^2 u$$
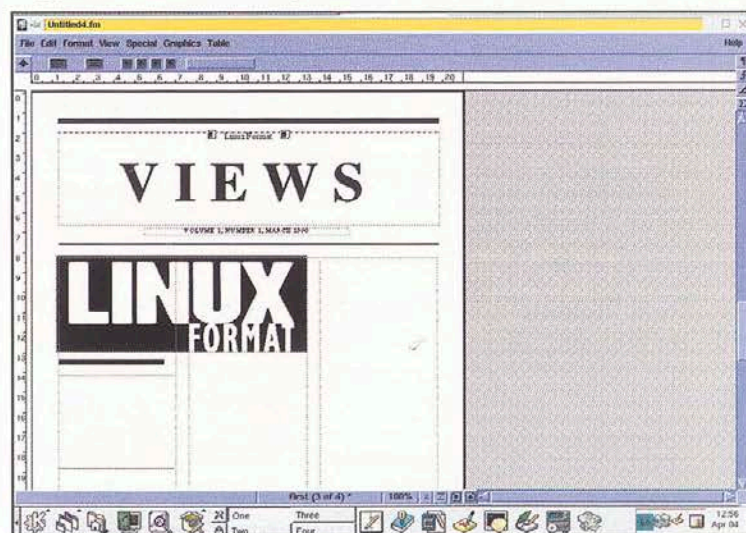
Multicolumn layouts are possible, as are conditional text elements and various graphic formats– set up the master pages properly and you're away!

Quite a few aspects of the software may appear daunting at first, even if you are experienced with DTP or layout software. Fortunately there is a handy guide to getting started, which is supplied, naturally enough, in the *FrameMaker* format.

This is still the beta version of the software, so don't expect it to run flawlessly. For example, on one of our test machines we had problems with graphic frames disappearing. That said, it does seem to be incredibly stable.

To use *FrameMaker* properly, you will need to sign up online for a licence number, which is then emailed to you. This licence is currently valid until the last day of the year, which should certainly give you long enough to try out the software. **LXF**

LINUX FORMAT Verdict

Certainly worth the download, but don't expect a fully fledged professional design system. *FrameMaker* is strongly oriented to structured documents such as reports and books, and in these areas it performs very well.

Rating looking good!

PROGRAMMING WITH QT

Publisher: O'Reilly **ISBN:** 1-56592-588-2
Author: Matthias Kalle Dalheimer **Price:** £21.95

For those unfamiliar with *Qt*, it is not really a programming language, but rather a well developed C++ class library. Its primary concern is in interface design, and it's probably most analogous with *Motif* on Unix systems. Another thing to note is that *Qt* code can also be recompiled to run on Windows NT and Windows 95/98 platforms, making it very useful for cross-platform development.

The aim of this book is quite simply to help you understand what *Qt* can do for you, and to guide you through the learning process. The *Qt* distribution does come with a fair amount of detailed documentation and tutorial material, but the author feels that the learning curve can be quite steep, and the official material is not very beginner friendly.

The book assumes a fair working knowledge of C++, which is reasonable, but starts off with the absolute basics of *Qt*. The chapters follow a logical progression and introduce components in an easy to understand way. The examples build up to include more widgets and classes, until you should be able to tackle any GUI-based application.

Each chapter ends with some example exercises which test your knowledge. In fact, the book could have done with more of these, as it is supposed to be teaching, rather than acting as a reference book. At 360 pages it isn't huge, and it could have offered more value by extending these sections, but the exercises included are well worked out.

Later chapters of the book also cover some more diverse topics. There is one on the use of GUI builders such as *QtArchitect*, *QtEZ* and the like. Another deals with *Qt's* application to

With easy to follow chapters, this book is a great way to get to grips with the QT library system.

network programming, and there are sections on writing plug-ins for Netscape, interfacing with Perl and an examination of some of the most successful *Qt* applications.
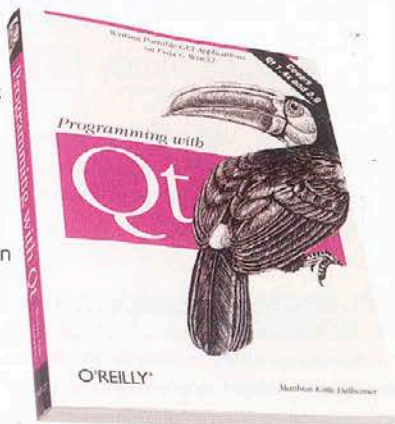
This book is an excellent alternative to wading through the official documentation and provides a reasonably easy way to learn this powerful library system. In spite of the inconvenient method of getting hold of example source code (downloading from a website), it is easy to follow, well written and full of information.

Prospective *Qt* programmers don't have a lot of choice at the moment, though with the release of *Qt 2.0*, there may be a few more titles on the shelves. For now though, we can heartily recommend this volume.

LINUX FORMAT Verdict

An excellent introduction to Qt programming, which is pitched just right for C++ programmers who want to take advantage of Qt's excellent features.

Rating 9/10

LINUX: The Complete Reference

Publisher: Osborne-McGraw Hill **ISBN:** 0-07-212164-5
Author: Richard Petersen **Price:** 24.99

This is the third edition of the *Complete Reference*, and since Richard Petersen has written them all, I think we can agree that he must have got it right by now. He teaches Unix courses at Berkeley and the University of California, so it's probably a well founded assumption.

This book is a complete reference, but if anything, it's biased towards beginners. It starts with an overview and concise history of the Linux OS and makes a good case for using the OS in the first place. The lengthy introduction section also covers various distributions, installing and basic configurations.

The early installation section dwells at length on Red Hat and Caldera's OpenLinux. This is certainly worthwhile because the book is accompanied by single-disk distributions of both – Red Hat 6.0 and Caldera's OpenLinux 2.2. Both are definitely worth having and are easy to set up for beginners, even though they are no longer the latest versions.

Further sections explain how to connect to, and make use of, the Internet, system and network administration, plus various selected applications. There is a huge section dealing with the two main window managers, *KDE* and *Gnome*, which is particularly helpful for newcomers as they are the most common defaults for most distributions.

A little more detail on some of the alternatives wouldn't have gone amiss though. As might be expected, there is also a large and very informative section on *Apache* and various other server options. It wouldn't be too much to say that armed with this book, the network-savvy reader could probably configure a Linux server

without too much trouble. The 18-page index is very thorough, so it is quite easy to use this book as a manual, although many items with multiple page references could do with clarification, especially as the book bills itself as a complete reference.

In fact, if there is one criticism to be levelled here, it would be that this book is slightly
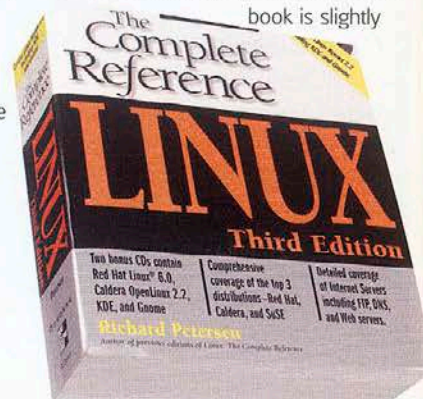
Despite the title, this book will be of much more use to the beginner than the experienced Linux user.

mis-named – it's definitely more of a beginners' reference guide. It would be impossible to cover the same range of topics in one volume to the depth required of a true, comprehensive reference, but Linux experts might be better off surrounding themselves with reference books on specific topics.

For beginners and intermediate users though, this is a very useful reference book indeed.

LINUX FORMAT Verdict

Experts may notice the lack of some detail on specific subjects – the book is really more suited to beginners. It is well written and well planned though.

Rating 7/10
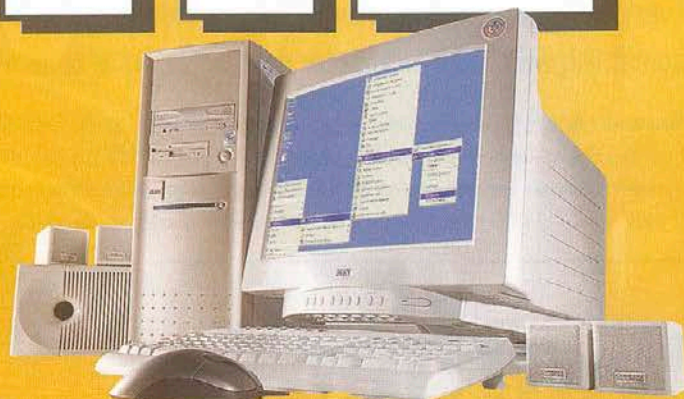
HOT PICKS

Every issue **Richard Drummond** will be rounding up the hottest new software that has come to our attention. If you can think of something we should include for next issue, email us at linuxformat@futurenet.co.uk

You won't find any ratings or verdicts in this section, because all the software here is some of the best available to mankind. These are our Hot Picks for the month – software we reckon is more than worth the effort of installing. You'll find a wide variety of stuff here, from mail clients to theme managers, databases to

fractal graphic generators – the only criteria is that it's hot software. If you have developed an application, or know of a new and really cool piece of software, why not write in and tell us about it. In fact, why not write a short piece about why you think it's so good, and we may include it here!

wmtheme

Function: Desktop configuration **Version:** 0.6.4

Website: wmtheme.snowblind.org **Packages:** deb, tgz

One of the great features of the latest round of window managers for X is 'theme-ability'. That is, their visual appearance can be customised to your personal taste with user-defined colour, widgets, borders and so on. If you can't be bothered to roll your own themes, fear not: dozens of prepackaged themes are available for

download from sites such as http://www.themes.org/. However, the process of finding what themes you want, downloading, unpacking and installing them can all be a bit too much of chore.

Not any more. *wmtheme* is a Perl script, designed to be run from the shell, which can simplify the

management of themes. It currently understands themes for many of the most popular window managers, including AfterStep, Blackbox, Enlightenment, Sawmill and Windowmaker.

Installing a theme is just a matter of calling the script with the name corresponding to your window manager, supplying the option --get and the name of the theme you are looking for. *Wmtheme* will log into the themes.org website and find all the those which match that name. For example, "etheme --get Blue_OS" would retrieve the BlueOS Enlightenment theme. You will be warned if a particular theme might be incompatible with your version of the window manager.

On the successful downloading of

your chosen theme, *wmtheme* will unpack it, install it in the appropriate place and activate it. If you decide that a theme doesn't suit you, you can easily get rid of it with *wmtheme*'s backstep option.

Wmtheme also provides other functions for the general management of themes. You may list the themes currently installed on your system, rename or uninstall them. It can activate an installed theme by name and set it as your default.

Wmtheme automates what was previously a rather tedious job. Now we can all have good-looking desktops (taste permitting) with the minimum of fuss. If the authors were to add support for GTK+ and KDE themes and build in a graphical interface, it would be nigh on perfect.
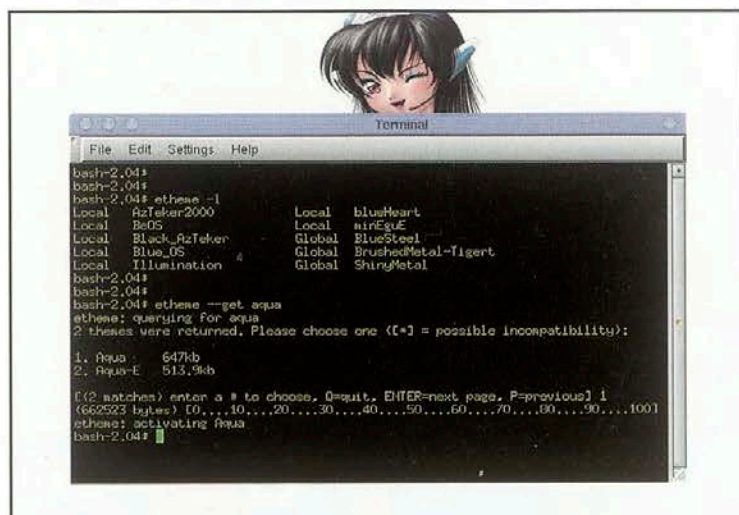
Whether your tastes lean to Manga, fractals, or something more wholesome, this theme manager will let you change your look with ease.

Grip

Function: CD player/ripper **Version:** 2.93

Website: http://www.nostatic.org/grip **Packages:** RPM

Whoever first thought of hooking up a CD player to a computer ought to be awarded a medal. Not only do you get buckets of cheap storage space, but –

as a bonus – you get to listen to regular audio CDs too. Your average CD-ROM drive lacks the physical buttons and controls of the hifi CD

player, but this can be more than made up for with a software replacement.

Grip provides a fairly standard graphical interface for playing audio CDs on your Linux box. It comes with the usual advantages of a software-based CD player. For example, *Grip* can maintain a database of CD titles and track listings, so that you'll never need to dig out the CD box to find out what

- song is playing. In common with most similar programs these days, *Grip* understands the CDDB protocol. This a method of saving you the hassle of keying-in all those track listings yourself. Whenever you insert a CD it doesn't recognise, *Grip* can automatically query a remote CD database server via the Internet and download the relevant information for that disc (the database is HUGE).

The most notable feature of *Grip*, though, is that it can extract or 'rip' audio tracks from a CD and save them to disk as samples. It can perform this either with its built-in ripper, which is based on the *cdparanoia* package, or by calling another external program to do the job.

Grip can optionally spool the track being extracted to an external MP3 encoder, in effect providing an easy-to-use frontend for the creation of MP3 tracks. *Grip* also helps to manage your digital music collection: it can automatically tag the MP3 tracks it generates and it also supports the author's DigitalDJ MP3 jukebox system.

Grip's GUI is created with GTK+ toolkit and, while rather staid in appearance when compared to some other CD players, it is clear to use. It has stacks of options for you to twiddle with and tune the environment for your listening pleasure. In addition, the program may also be controlled via the command line, a method particularly suited to the batch encoding of discs.

The purpose of *Grip* is not only to provide a comfortable means of playing your audio CDs but also to facilitate the creation of a 'computerised' copy of your CD collection. At both roles it succeeds.

Grip is not just a substitute for your stereos CD player...
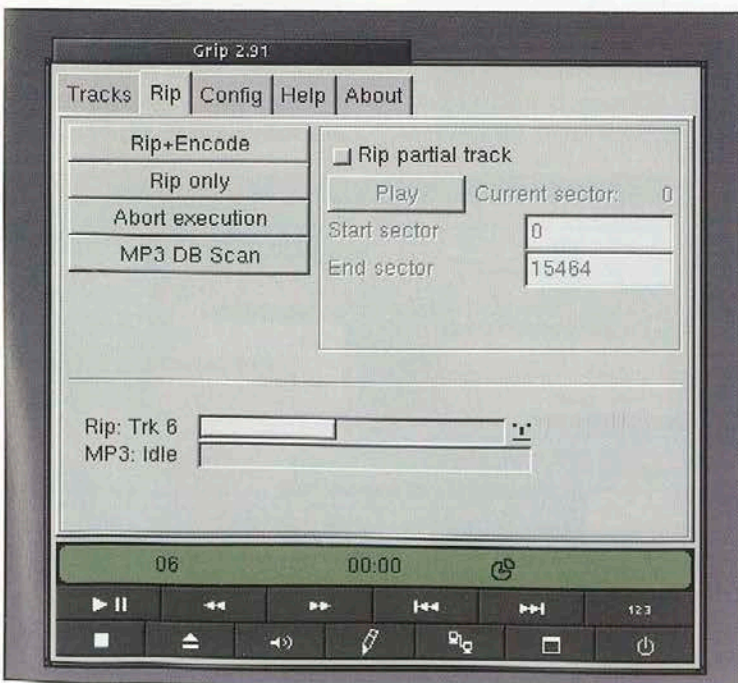
...you can throw away that old tape recorder, too.

Mahogany

Function: Email client **Version:** 0.50 **Packages:** deb, RPM
Website: http://www.wxwindows.org/Mahogany

Time for a confession. I have yet to find a graphical mail client for Linux that compares with *YAM2.0* on the Amiga in terms of reliability, power and good GUI design. Don't laugh. *YAM* knocks the spots off *LookOut Express*, *Netscrape Mail* and anything else you'd care to mention. I'm still looking, though, and that is why I was interested when I came across *Mahogany*.

Mahogany is a combined mail and news reader based on the cross-platform GUI kit, wxWindows. With all the attention paid to GTK+ and QT, you may not of heard of wxWindows, but in many respects it is a much more mature system. It currently supports the usual varieties of Windows, MacOS and various Unices and on Unix platforms can be built to use either the Motif or GTK+ widget sets. The latter is the norm on Linux. Mahogany ships as binaries for Windows, Solaris and Linux with a MacOS version in the works.

Mail in *Mahogany* is organised into a tree of folders. Folders can either be a mail (or news) source - such as your local mail spool, a remote POP3 or IMAP mailbox or an NNTP server - or a

filesystem object in which mail is stored. Folders may be stored in the normal mbox format or MH format. Mail is picked up from any of the sources and deposited in the InBox folder where it can then be filtered to other folders as desired.

This hierarchy of folders provides *Mahogany* with a lot of its power. The properties of each folder may be configured individually or it can inherit the properties from its parent folder. You can change everything from the account identity used to log in to a mailbox, to the way mail is viewed in that folder, to the signature and reply headers used when composing mail. Luckily, *Mahogany* has beginner and advanced modes so as not to bewilder the novice with too many options.

Other strong features of *Mahogany* are its built-in Python interpreter and extensible plug-in system. Plug-ins are provided to use Python script for mail filtering and another for synchronisation with a Palm PDA. Mahogany is fully MIME-compliant and can display images inline with viewed mail. Since it doesn't rely on a particular desktop →

→ for its MIME handling preferences, it manages these settings itself. You may specify external 'helper' programs to handle particular types of attachment, for example, a browser for viewing any HTML attachments. *Mahogany* also has extensive address-book facilities built-in.

In spite of all these good things, *Mahogany* still needs a lot of work,

especially in its GUI. It's main window uses the now familiar three-pane view, showing the tree of folders, the currently open folder's contents and the current mail being read.

That's fine. But I would like to be able to control all these three elements simultaneously from the keyboard. It's a real pain when you are dealing with a lot of mail having to resort to using the mouse to move to the next mail or jump to another folder. More keyboard shortcuts are definitely needed.

And then there's the folder tree itself. I want to be able to arbitrarily order each level of the tree myself rather than having them sorted alphabetically. My final complaint is

that *Mahogany* is woefully slow at handling its mail folders. Is just too frustrating having to wait while *Mahogany* retrieves new mail, filters mail or opens a new folder. With large folders containing many mails it just grinds to a halt.

After trying out *Mahogany*, I have reluctantly reverted back to *KMail* for my mailing needs, even though *Mahogany* is much closer to fulfiling my functional requirements.

Why? I just found it too slow and cumbersome to deal with the large volumes of mail I regularly receive. However, I will watching its development very closely in the future – if the developers can just sort out these few problems, it'll be great.
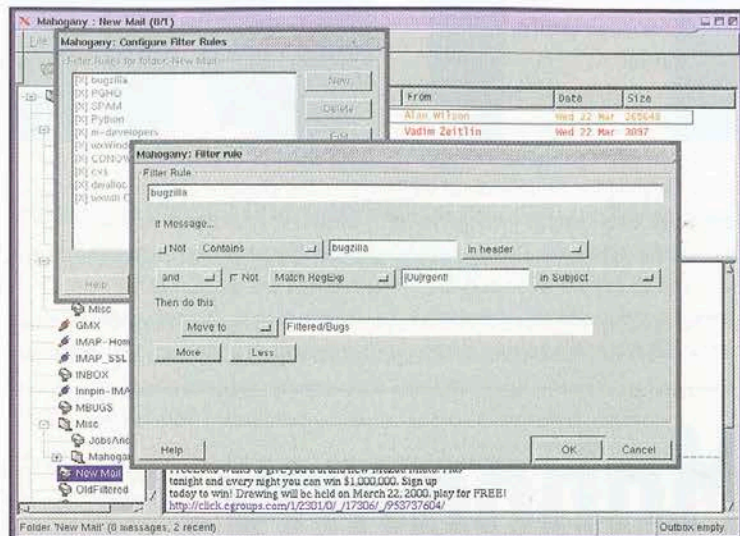
Mahogany's powerful filtering capability is mainly thanks to its built-in Python interpreter.
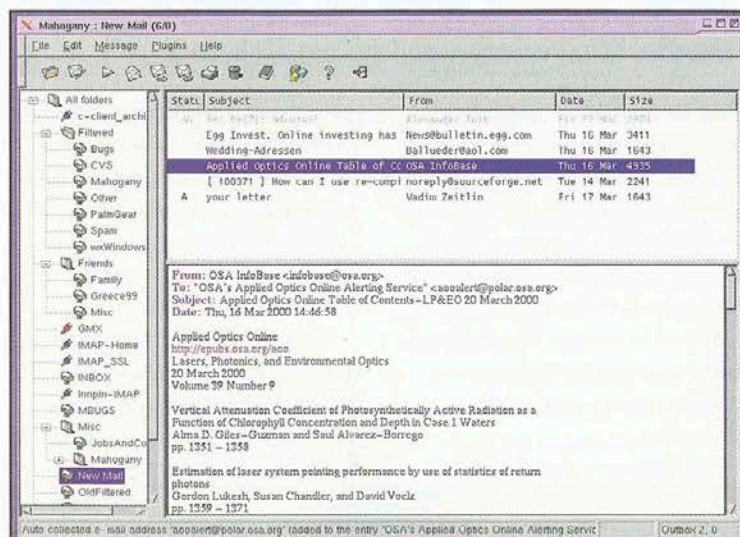
Mahogany in action. Does every mail reader really have to named after a type of tree? I've got dibs on Giant Redwood...

THE COMPETITION

Choosing a mail client is very much a matter of personal taste. Why not check out these alternatives?

KMail (http://www.kde.org)A QT-based mail-client with a great GUI, but rather lacking in features. It will work outside the KDE desktop, but not at its best.

Spruce (http://www.xtorshun.org/spruce)Another mailer with a well-designed GUI, this time using GTK+, but again rather basic.

Balsa (http://www.balsa.net)A good-looking but very incomplete mail reader built for the GNOME desktop.

Postillion (http://www.postillion.org/)A capable mail client but is beginning to look rather dated.

Gaby

Function: database **Version:** 1.9.19 **Website:** http://gaby.netpedia.net/ **Packages:** source only

As befits a UNIX-like operating, Linux is well catered-for with professional database applications software. But running a full-on SQL server just to manage your collection of video recordings is bit much. *Gaby* is a GTK-based personal database manager designed expressly for more modest tasks. It requires various shared libraries

from GNOME, but doesn't actually need the GNOME desktop to run.

A *Gaby* database consists of a series of tables. A table corresponds to what would be a relation in a relational database system and consists of a number of fields. Each

Creating a simple database is easy with Gaby's graphical builder.
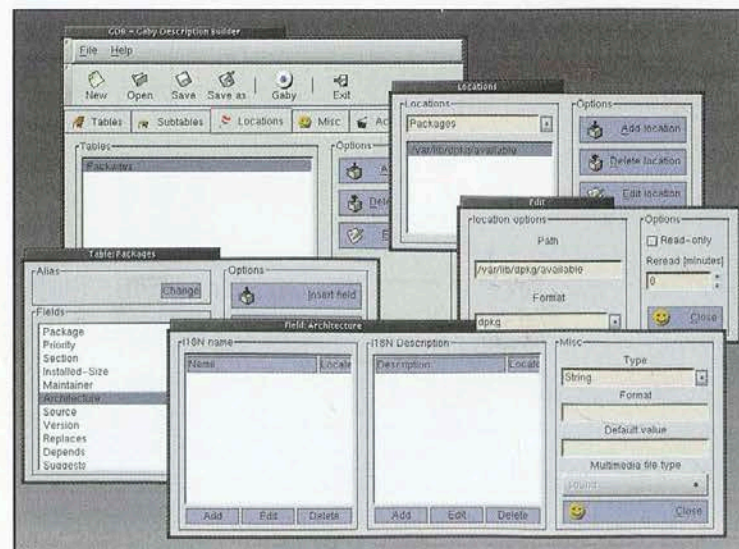
table can be associated with a particular file - which *Gaby* calls a location - and may be stored in variety of formats. This is useful for exchanging data with other applications. Field types supported include the usual string, numeric, boolean and dates.

Multimedia data types can also be handled. If you have *imlib* installed, you can use images as fields, while, if you are running *EsounD*, you can use audio samples. Each table or can be displayed on screen in a variety of ways, as a list, a form or in a more specialized format. You can produced 'reduced' views of a table by creating subtables, picking which of its fields you want displayed.

The main *Gaby* executable is used to manage a *Gaby* database. Depending on how the database is configured, you can view, sort, search the database and enters new data. The other half of *Gaby*, is to actually create a new database. A database scheme is described in GabyScript, a flexible scripting language designed for the purpose, and is stored in a description file.

A description not only describes the structure of a database, it also specifies it's visual appearance and behaviour too. Descriptions can either be created and edited by hand, or you can use builder, a point-and-click

database description editor. This is not complete yet, so to take advantage of some of *Gaby*'s more advanced features like actions, you still need to know GabyScript.

Action are best illustrated with the example databases supplied with *Gaby*. The default application is an address book and manages a list of contact details, names, address, phone numbers and email addresses, etc. Actions supported on this database include dialing a phone number and sending an email. These actions are implemented in GabyScript and make use of *Gaby*'s net plug-in, a module which provides functions for interfacing with Internet applications. Alternatively, an action can be a Python or shell script. Another example of an action is with the supplied CD database which has the ability to import details from a CDDB server.

If you don't need to create your own database from scratch, *Gaby* is incredibly easy to use, deceptively so in fact. The end user of database need never know of the power that lies hidden under the bonnet.

Specifiying your own database or tweaking an existing one is also simple thanks to the builder tool. Getting to grips with *Gaby*'s more advanced functions is more difficult, however. This is not due to any inherent complexity, but because of a

dire lack of adequate documentation. The numerous examples should point you in the right direction, though.

This is certainly one of the most manageable and flexible database aps we have come across.

SOCIABLE SOUND

As mentioned above, *Gaby* uses *EsounD*, the Enlightened Sound Daemon, for playing audio samples. What's that then? *EsounD* is a solution to the problem of how several programs can generate sound simultaneously on a machine with a single soundcard.

EsounD mixes the multiple sound streams generated by whatever programs are running and outputs them as a single stream to your audio hardware. It is no longer necessary for a program to hog the sound device all to itself. Other benefits include network transparency: you can play a sound on one box and listen to it on another. See http://www.tux.org/~ricdude for more information.

FreeAMP

Function: MP3 player **Version:** 2.0.6 **Website:**
http://www.freeamp.org/ **Packages:** deb, rpm, slackware

MP3 - or to give it it's full name, MPEG I or II Layer 3 – is quickly becoming the format of choice, both for music storage and the transmission of audio data across the Internet. This is because it offers a good degree of compression without sacrificing sound quality. To make the best of the MP3 format on your Linux machine, you need a comfortable and feature-rich

MP3 player. Enter *FreeAMP*.

FreeAMP is an open-source MP3 player which shares much in common with the WinAMP player for Windows. Current versions exist for Windows and GNU/Linux and ports are currently in progress for the Mac, BeOS and Solaris operating systems. *FreeAMP* provides a familiar, CD player-like interface whose visual appearance is →

Don't be left behind in the music revolution. Get with it and get FreeAMP.
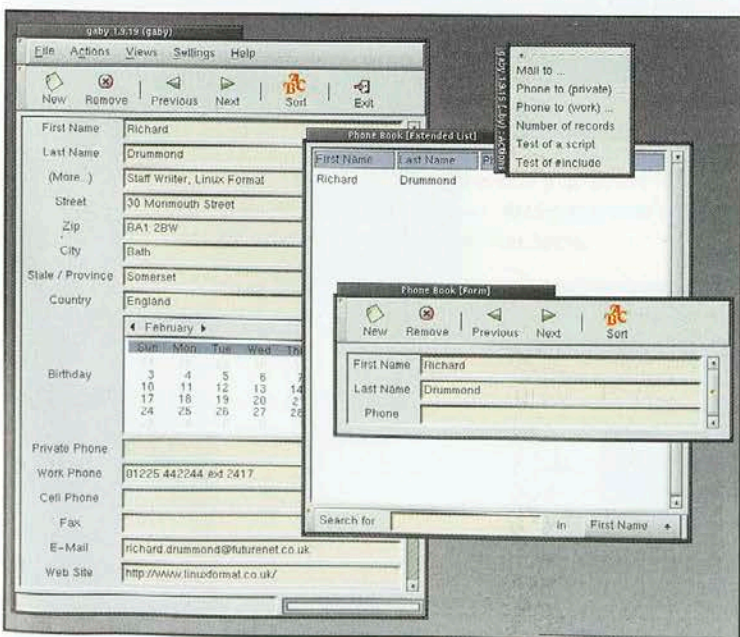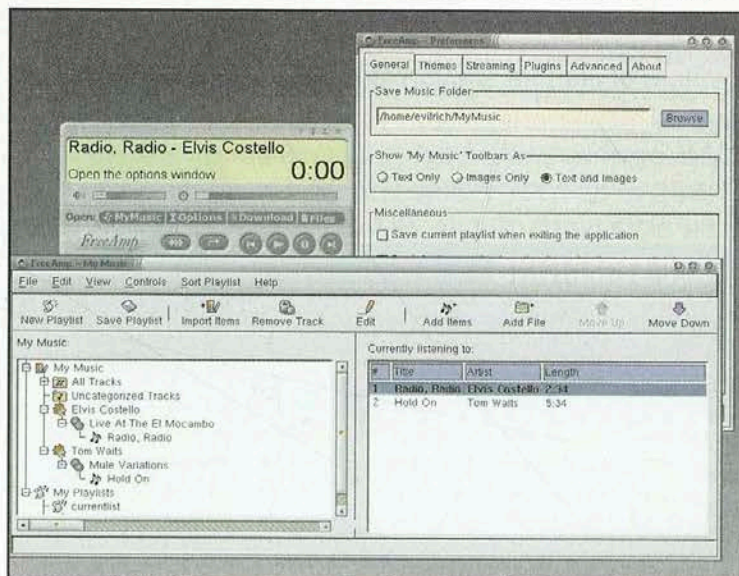
The default Gaby database is a address book. Not very exciting, but useful.
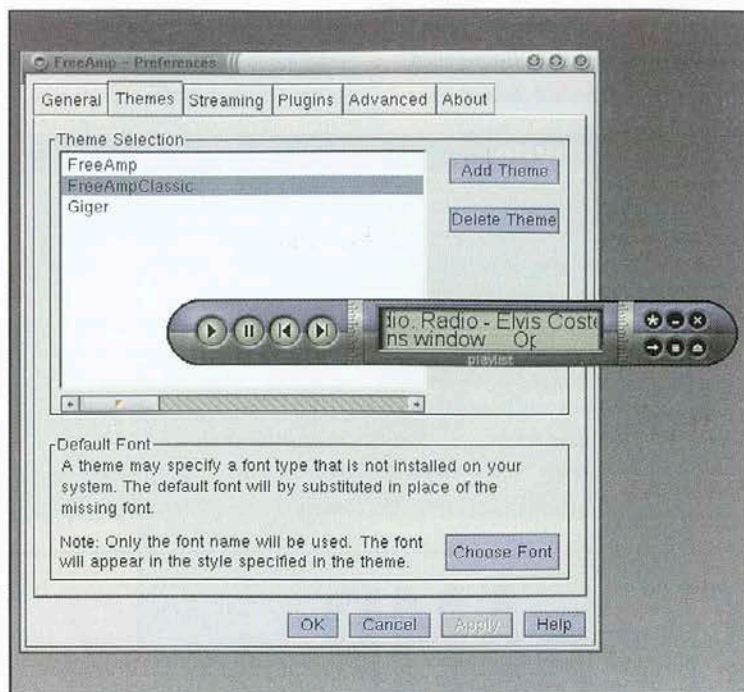
→ customizable with 'skins'. It can either play tracks that are stored locally on disk or can stream live audio from the Internet, supporting icecast and ShoutCAST servers.

On Linux, *FreeAMP* can generate audio via the Enlightened Sound Daemon and supports the modular sound system, ALSA (Advanced Linux Sound Architecture). The MPEG decoder employed is a tweaked and optimized version of the GPLed *Xing* decoder.

FreeAMP supplies a host of tools to help you organise your music collection. The *MyMusic* browser can automatically scan your hard drive to find what audio tracks are locally available and it groups them according to artist name, album title and track name. You can also categorise and sort tracks by different genres if you like. Tracks can be added to a playlist simply by drag'n'drop. Getting tracks from the Internet is a breeze thanks to

You can change the look of FreeAMP to suit your mood or the music you're listening to.

FreeAMP's built-in download manager which understands Real Networks's RealJukeProcess download process, a system used by a large number of MP3 portal sites.

Currently the *MyMusic* browser can only handle tracks stored locally on your hard drive, but in the future it will be able to understand tracks stored on Internet streaming servers and regular audio CDs. This should provide a seamless environment for listening to music no matter the format or location.

If you want to get in on the MP3 revolution, then *FreeAMP* is a great tool to do it with. It is well-designed, easy-to-use, has comprehensive documentation and good playback quality. Need I say more?

FREE STREAMING

Streaming audio means that an audio track is played as it is being downloaded from the Internet – rather than having to wait for the complete track to be transmitted to a local file and then played from there. *FreeAMP* supports two streaming formats, SHOUTcast and icecast, both based around MP3 encoding. SHOUTcast employs proprietary server technology but the server may be downloaded for free for non-commercial usage. icecast, on the other hand, is fully GPLed, so its use is unrestricted. See http://www.shoutcast.com/ and http://www.icecast.org/ for more information.

Terraform

Function: Fractal landscape generator **Version:** 0.6.6
Website: http://terraform.sourceforge.net/ **Packages:** RPM

Terraform is described by the authors as an "open source interactive height field generation and manipulation program". To the layman it's a fractal landscape generator with a swish GTK+ interface.

Terraform calculates a landscape from a random number. This landscape is displayed on screen and can be manipulated with various mathematical transformations to achieve the vista you are looking for. Some operations have names which at least suggest the effect they might have: Craters, Guassian Hill, Errode and Rough&Smooth; others are more generic, such as Transform.

The key to success here is experimentation (the documentation is pretty poor anyway). Have a play around with the different tools and get a feel for what they do. Most of the operations have a realtime preview window in their dialogues, so that you can see how altering the parameters will change the elevation of your

landscape. The lack of a history and an Undo function is an annoyance, but you can easily clone a landscape and work with a copy instead.

Loading and saving of landscapes is performed as elevation maps in the Targa bitmap format, but other formats are supported. The colour of each pixel corresponds to the height at that point of the map. This format isn't terribly useful, but you can have some fun by loading real images into *Terraform*.

More interesting is the ability to export the landscape as a 3D model to be rendered by the POV ray tracer. In future it will support OpenGL rendering as well.

At the moment Terraform is more of a toy than a useful tool. It's instability and lack of documentation make it a real pain to use. Things could change with some more work, however.

Landscape gardening on a grand scale with Terraform. Where did I put my shovel?

Moonshine

Function: IDE **Version:** 1.0.2

Website: www.suite3220.com/moonshine **Packages:** RPM

Program development on UNIX-like platforms has been traditionally been done with shell tools, command-line driven compilers, configure and make scripts. This can be rather off-putting to a developer who is used to the graphical IDEs and Visual-type environments available on other platforms. *Moonshine* is a package which attempts to redress the balance.

Moonshine is a language-independent graphical IDE for X, built with QT. It is shipped with modules geared for C++, QT, QT2 and web development. These modules are just configuration files which describe the look and feel of *Moonshine* and can readily be adjusted for your needs, whatever language you are developing in. The documentation describes this in some detail, and the customisation process isn't too taxing.

The main *Moonshine* window is the project browser and provides a tree list of the components of your project. Context-sensitive menus launch actions according to the filetype of a component. For example, Text files may be edited with the built-in editor, *Brew*. This is a basic but nimble programmer's editor with syntax highlighting and a macro facility.

A neat feature of *Moonshine* is its Capture facility. That is, the output of any shell commands it runs can be redirected to display in a *Moonshine* window – which supports syntax highlighting, cut'n'paste and a review buffer. *Moonshine* also keeps a history of all the commands it executes.

Making your own module or modifying one of the supplied modules is fairly easy and a necessary task to tune the environment to your own specific needs. A module specifies all the filetypes that the browser understands, the icons associated with each type and the various actions which may be applied to files of that type.

Actions can be any normal shell command or a built-in function. For example, a Makefile would support the calling of *make* as an action, while an image might support the launching of particular paint package for editing. Modules also configure the syntax-highlighting in Brew and the capture window.

Moonshine is a well-thought out package with plenty of scope for improvement (a drag'n'drop Makefile would be nice). It is ideal for avoiding culture shock when beginning to program on Linux.

Linux can have a graphical IDE, too. It's not all shell commands and switches, you know.

DISTRIBUTION ROUNDUP

Distributions are designed to make using Linux easier, but choosing the right one isn't always so simple. **Mike Saunders** looks at your choices...

Back in the early days of Linux, getting the OS onto your machine was a long and complicated task. Individual compressed packages had to be downloaded off the net and put together on your hard disk, with lots of tweaking of configuration files required before you had a fully-working system.

Enter 'distributions' – a concept designed to make installation and configuration of Linux much easier. Instead of working out which software you needed and then trying to download it, everything required was supplied in a simple package, along with manuals and extra software CDs.

Free for all

There are many Linux distributions available today, with more appearing all the time. Because Linux and its source code is freely available, anyone can grab the bits and pieces of a working system and put them together to make their own distribution. Clearly Red Hat is the most prominent distribution vendor, but other companies like SuSE and Mandrakesoft are following closely. This freedom for anyone to make a distribution has led to an assortment of packages with different aims – some try to provide a fast and easy installation, some head for being reliable servers on networks and some just want to be all-rounders.

This range of focus looks set to expand as Linux becomes more popular in other markets. You can buy full boxed-sets with manuals, extra discs containing applications and the source code, plus invaluable installation support. If you're familiar with Linux or just want to upgrade, you can get hold of a cheap single CD version of the distribution, or download it off the Web

Choices, choices...

So how do you choose the best distribution? It really depends on what you're looking for. If you're new to Linux and want a simple introduction, it's best to find one with a friendly graphical installer and plenty of software. This way you'll be able to get up and running with minimal fuss, and won't have to trawl the web trying to find decent programs. On the other hand, experienced users will value reliability and flexibility over flashy installers and configuration tools. If you need a distribution to run on a small file or print server, the latest dazzling desktop environments aren't going to be of any interest.

One of the most significant obstacles for newcomers installing Linux is the partitioning process. If your whole hard disk is devoted to Windows, making enough space for Linux is a complicated and daunting task. Distribution engineers have devised various methods for dealing with this, ranging from automatic

COMPATIBILITY

Many distributions are based on others, which is important to note when you're trying to find compatible software. Because the majority of software in a typical distribution is free and open source, another company can take the distribution's framework, add their own tweaks and improvements and market it themselves. One example of this is Linux-Mandrake, which began as Red Hat with KDE ready-to-go (Red Hat had been avoiding the desktop suite because of licence issues). After having a good, solid base to work on, however, Mandrake has become a worthy distribution in its own right.

partitioning software to placing the whole installation inside a directory on the Windows drive.

Desktops

The choice of default desktop is also a crucial one. While experienced users will have no trouble in setting up their favourite window manager or desktop environment, newcomers will want things ready for action straight away, without having to fiddle with configuration files. *KDE* and *GNOME* are the biggest contenders for the desktop, providing tools, utilities and a drag-and-drop environment, but simpler window managers like *Window Maker*, *Sawmill* and *IceWM* are often better alternatives on systems with smaller memory. In this test, we've also taken note of how well the desktop is set up. Has it just been thrown in, or have the developers configured it to work with the special features of the distribution?

Version number is also a key issue to watch out for. While it appears good to have the latest releases of popular packages, they're not always the most stable. New features need heavy
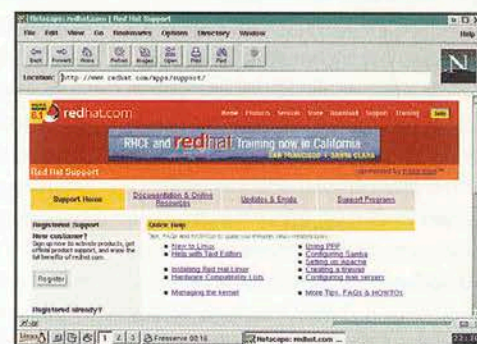
Detailed online help files are absolutely essential when you're getting started.

testing, and it's often better to have an older, tried-and-tested version than the current development release. This is especially true for production machines and servers – on your home PC you can gamble with newer, untested packages, but you might not want to use your business as a testing zone. In general, most distributions stick with stable versions of important programs, like the kernel and related tools, and occasionally have in-development versions of less critical applications.

The contenders

Here we take a look at eight of the major distributions, examining the areas discussed above and determining the type of user each distribution would be suited to. We're looking at the single CD versions here, which are available from many retailers in the UK. Most of the CDs in this roundup were kindly supplied by the Linux Emporium, who carry a wide range of both single CD and boxed distributions. You'll find their website at **http://www.linuxemporium.co.uk**

Corel Linux

Corel's entry into the Linux market certainly has a lot of potential, but there are problems too...

Corel Linux 1.0 – http://linux.corel.com

When Corel announced that they would move into the Linux distribution market, the news was greeted with both hope and apprehension by the community. While it was clear they could help to improve Linux's position on the desktop, many thought they were in danger of dumbing-down the OS. The software company, responsible for *WordPerfect* and *CorelDRAW*, have based their package on the respected Debian distribution, and aimed it at new users who have switched from Windows.

Corel's graphical installer is welcome news to new Linux users and those who have switched from Windows, featuring mouse and dialogue-driven configuration. As well as taking over the whole hard drive, you can install directly into a Windows directory. However, our system froze just after it had copied the files onto our drive, and reports from other users in the main Linux newsgroups confirm that this is a common problem.

According to the detailed support pages on Corel's site, these difficulties are due to hardware probing. Fortunately, though, the files were already in place on our drive so we could still start Corel Linux up.

> ## Our system froze just after it had copied the files onto our drive...
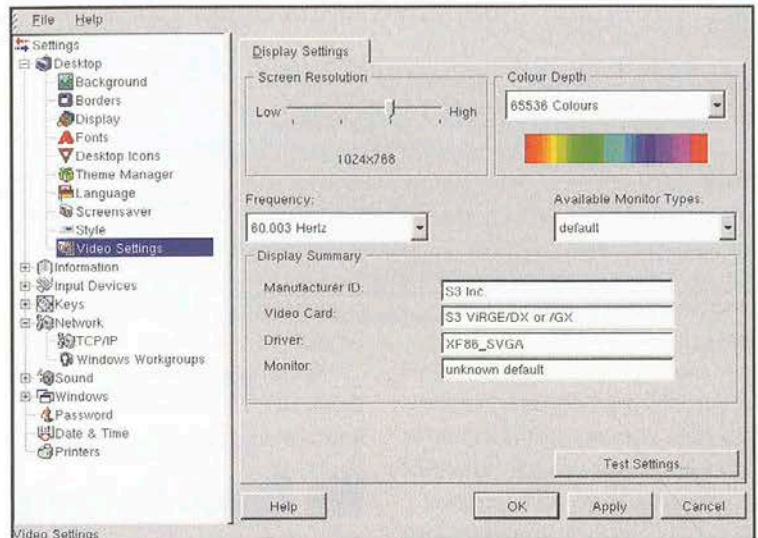
Network and video configuration can be done from the control panel.

Desktop details

Corel's desktop is based on *KDE* with various tweaks and enhancements. The most significant addition is the *Corel File Manager*, which resembles *Windows Explorer* and provides a familiar interface for new users. *FVWM 95* is also included as an alternative. The *XFree86* version is 3.3.5, and also supplied is the *GIMP* 1.0.2, *Netscape* 4.7 and Adobe *Acrobat Reader.* This is a reasonable range of desktop applications, although adding extra software is trickier as Corel Linux uses .deb packages, rather than RPMs (which most of the other distributions use). Still, there are

many Debian packages around and the supplied *Alien* tool can convert RPMs to the native .deb format. For more advanced users, the gcc compiler and *Midnight Commander* file manager are also available.

Problems...

This is a promising distribution, and Corel appear to have worked hard in many key areas. While there's little variety in the available desktops, the ability to perform *X* and network configuration through the *KDE* control panel is a bonus. Advanced users may feel restricted though, and the sheer amount of installation problems we've come across and heard about makes it hard for us to recommend it.

If they can sort out these troubles and produce a reliable installer in their next release, Corel Linux will be worth keeping an eye on.

Corel's desktop is based on KDE, with various extra enhancements.

LINUX FORMAT Verdict	
Ease of use:	6/10
Documentation:	8/10
Features:	8/10
Software range:	6/10

Has potential, but needs work on the troublesome installer.

Rating 6/10
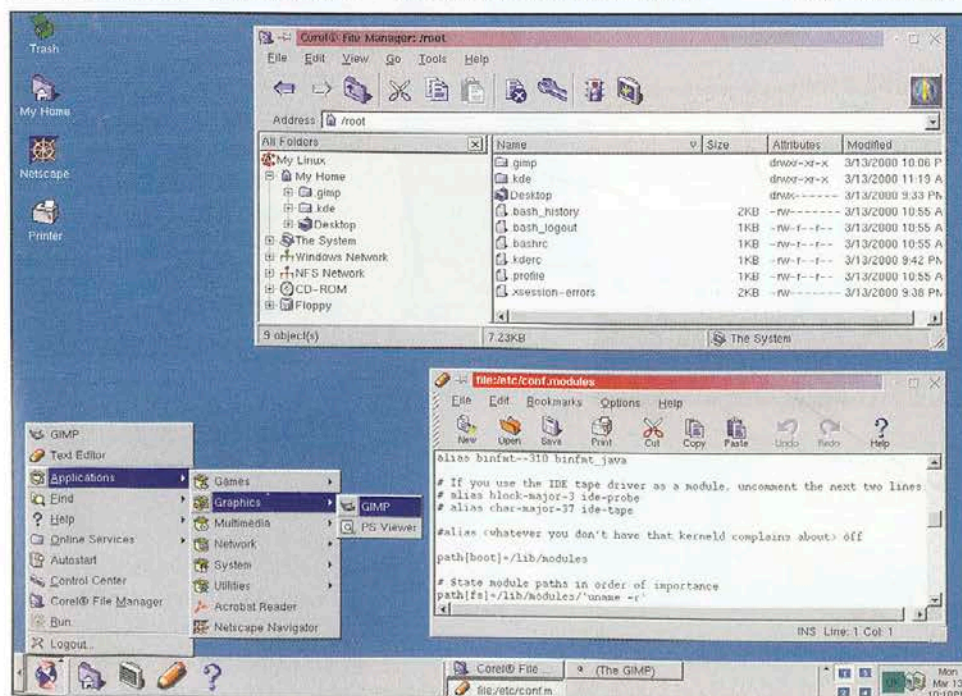
Definite

This UK-based company's Linux distribution is one of the best all-rounders we've seen so far.

Definite Linux 7.0 – http://www.definitesoftware.com

*The GNOME **desktop** with the Enlightenment **window manager**.*

Produced by Definite Software PLC, a UK-based company, Definite Linux is based on the Red Hat distribution. It aims to be up-to-date, with extra enhancements over Red Hat, and includes full UK support. The text-mode installation program is, like Red Hat 6.0, fast and straightforward. We chose our language and time zone, then used the simple *Disk Druid* tool to create our partitions. It installed our selected packages and detected our mouse and video card. After installing the boot loader and configuring a few extras, we were ready to go.
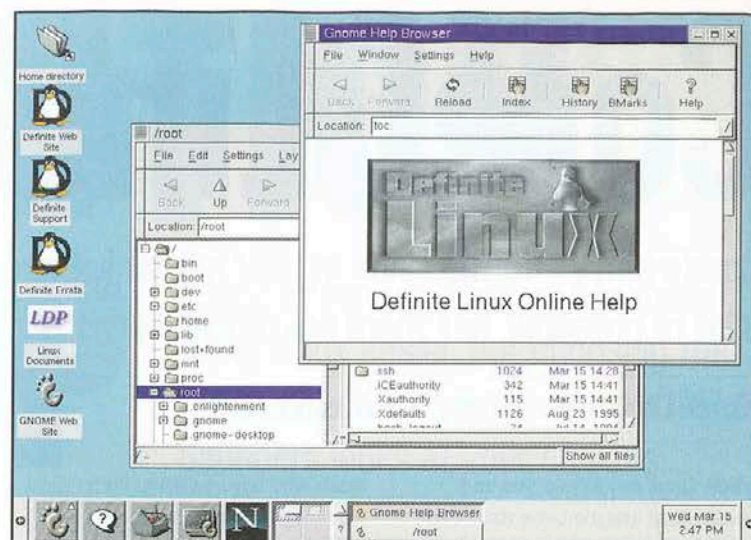
From starting the installation to booting our new Linux system, only 30 minutes had passed – very impressive,

and essential when several machines need to be set up. While many distributions are moving towards graphical installers, the speed of Definite's setup program is hard to beat.

GNOME, the default desktop with Definite, has not been altered much. Ready-prepared icons for the CD-ROM

We found Definite Linux to be an excellent and solid all-rounder.

and floppy drives would have been a nice touch, although typical users of this distro would be familiar with

configuring such things. Definite have added links to their support pages on the desktop, with extra links to other information sources in the documentation. *KDE* is also included for those who prefer it, and other window managers like *Window Maker* and *FVWM 95* are available too. In terms of applications, *GIMP*

1.0.4 and *Netscape 4.61* are joined by *Apache 1.3.9* and the *egcs* compiler. The supplied kernel is version 2.2.12, and the *XFree86* release is 3.3.5.

Plus points

Definite really excels under the hood. With RAID support, it makes a very capable server solution and not just a solid desktop OS. Another neat feature is the ISDN support, demonstrating the flexibility Definite provides. Also, this distribution offers improved security over its rivals. For example, *Netscape* supports full 128-bit encryption.

We found Definite Linux to be an excellent and solid all-rounder. The fast installation and decent range of software included makes it a powerful load-and-go solution, and the usual Red Hat configuration tools, together with *linuxconf*, create an easily maintainable system. Documentation and support are superb, helped considerably by Definite PLC being based in the UK. New users may find Mandrake easier to work with initially, but if you're looking for Red Hat with an extra bang for your money, this is it.
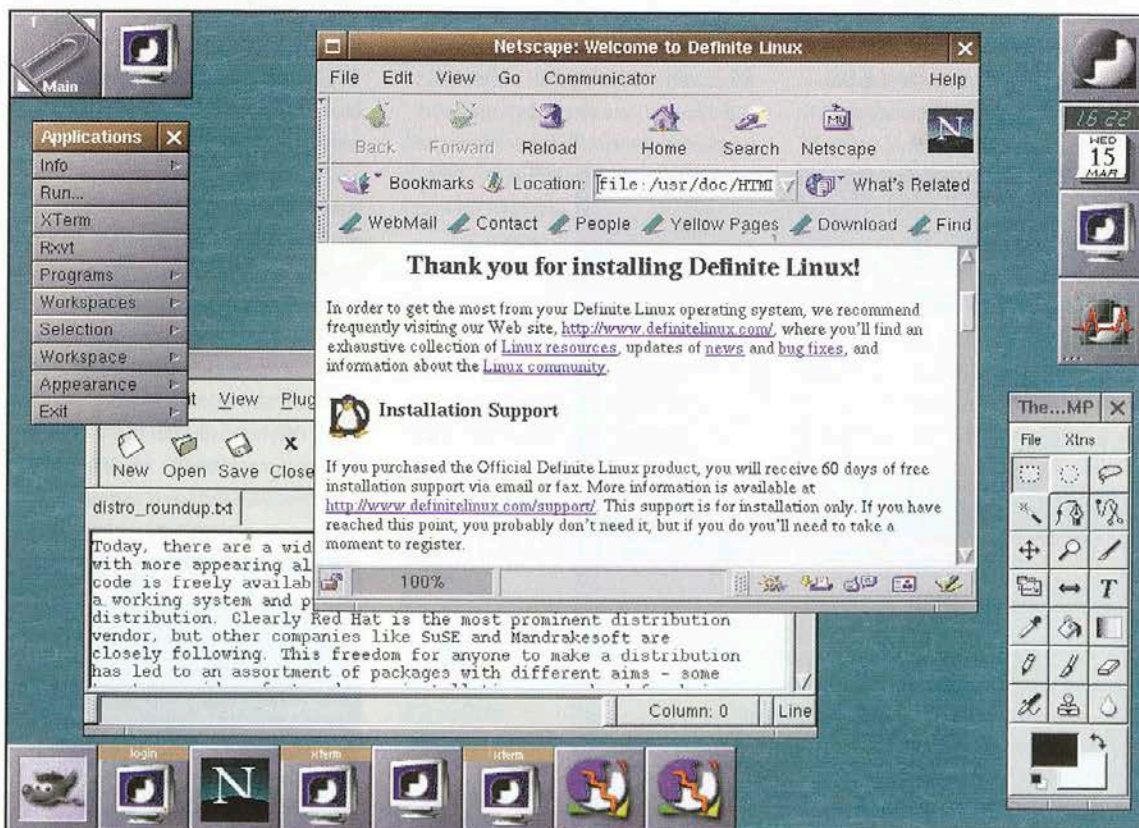
*Window Maker **is also supplied as an attractive alternative**.*

LINUX FORMAT Verdict

Ease of use:	8/10
Documentation:	8/10
Features:	9/10
Software range:	9/10

Powerful and flexible with some impressive features – a great choice all round.

Rating 9/10

WinLinux

The easiest to install Linux system in the world? After tests, we're not quite so sure...

WinLinux 2000 – http://www.winlinux.net

WinLinux 2000's primary selling point is that it can be installed onto a Windows drive without the need for any tricky partitioning. It's based on Slackware, while claiming that it is "the easiest to install Linux system in the Windows world".

By working with the *UMSDOS* file system, WinLinux uses a Windows directory to hold the Linux files, although it's slower than a proper, dedicated partition. We booted Windows 95 and ran the SETUP.EXE program. A typical installation utility appeared, and after choosing the 'Typical' choice for packages, the installer started copying files onto our C: drive. So far, so good.
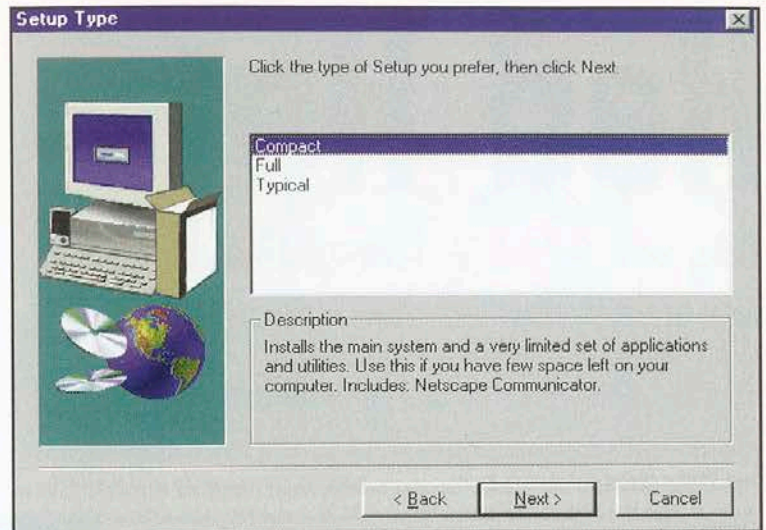
The installer finished and launched the configuration tool... and an error appeared. After dismissing it, we were faced with a blank configuration window. Pressing the Back and Next buttons only resulted in more errors, leaving us with Cancel as the only option.

Sadly, this wasn't just a problem with our system – we found many people complaining of the same difficulties on message boards all around the Internet.

Installation information

At WinLinux's site, we searched the Updates and Support section with no luck. We sent an email to their support team, hoping to find a solution, but had no reply. Eventually we found a link in one of the FAQ sections to an updated configuration tool. Why this was hidden away and not highlighted in the Updates and Support section is a mystery, but we finally installed the new configuration utility and it ran

The Windows-based installer offers a selection of installation sizes.

The installer finished and launched the configuration tool... and an error appeared.

better, although it only detected our mouse and not the video or sound card. When starting Linux, we also encountered error messages with some of the kernel modules.

WinLinux includes kernel 2.2.13, *XFree86 3.3.5* and version 2.1 of the glibc library, so it's completely up to date. The default desktop is *KDE 1.1.2*, which has been set up reasonably

KDE, the default desktop, has been tuned to provide extra space on-screen.

well, and the developers have changed the panel's size to fit smaller screens. *FVWM* is also available as an alternative system.

Other software which has been provided includes the *GIMP 1.0.4* and *Netscape Communicator 4.7*, while the *egcs* compiler is available for development work. This is a satisfactory setup, and *KDE* provides plenty of extra utilities for multimedia and the Internet.

WinLinux in action

In use, WinLinux's Slackware base makes it a solid and powerful system. Slackware's configuration tool (which is invoked with 'setup') is also a comprehensive utility, although no mention of it is given in the documentation. This isn't particularly surprising as the documentation is very small and badly-written.

WinLinux is clearly aimed at newcomers, but until they sort out the setup problems and provide some more detailed help, with some much clearer help files, we would have to recommend that you avoid it.

LINUX FORMAT Verdict

Ease of use:	4/10
Documentation:	2/10
Features:	6/10
Software range:	9/10

Good concept, let down by a problematic installation and poor documentation.

Rating 4/10

Red Hat

We see how these experienced distributors fare against the many new kids on the block.

Red Hat Linux 6.1 – http://www.redhat.com

One of the older distributors in this roundup, Red Hat have a reputation for making solid packages for both the desktop and server. The company created the popular RPM package format and remain the dominant figures in an increasingly competitive market. Like many others, they've made the move to a graphical installation process with *Anaconda*, which can also install in text mode.

While not as visually impressive as some of the other offerings, Red Hat's installer is clean, fast and presents useful help text at the side of the screen. It detected our mouse and

GNOME and Enlightenment, the default desktop for Red Hat.

graphics card correctly, and the *Disk Druid* partitioner was simple enough to operate. You're offered a choice of installation classes, from a typical *GNOME* and *KDE* workstation to a server, with a 'custom' option for more advanced users.

The installer also creates a normal user account, which is essential for a secure and reliable system.

Bundled software

Red Hat has put significant backing into the development of *GNOME*, which remains the standard desktop environment for its distribution. While debates over *GNOME* vs *KDE* are likely to continue for some time, it's still a
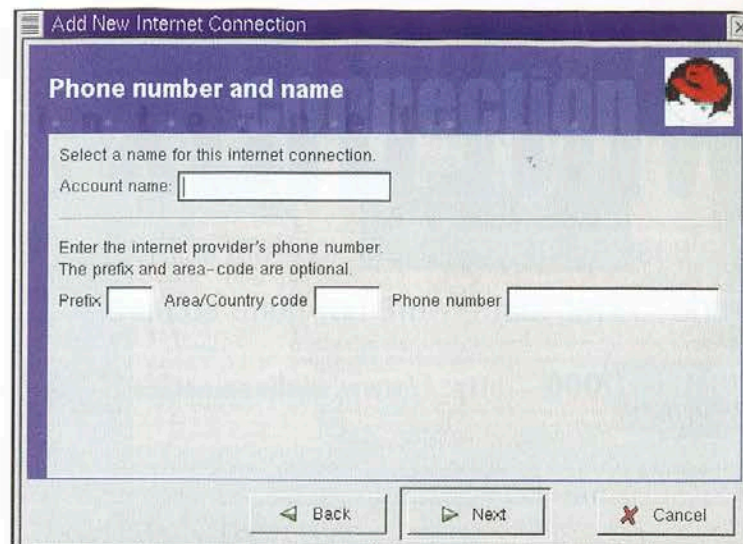
good choice and the number of programs and small utilities provided make a very rich and usable desktop.

Icons are ready prepared to point to the CD-ROM and floppy drives, while the *GNOME* help browser offers links to the Red Hat Getting Started guide.

The CD is also equipped with alternative window managers like *FVWM*, *Window Maker* and *AfterStep*. A good assortment of typical programs is supplied, with development tools, server applications (including *Apache* and *Samba*) and Internet utilities like *Netscape*, *Pine* and *Lynx*.

RP3 being used to create a new Internet connection.

Red Hat uses the *linuxconf* configuration program for common system administration tasks. While this is undoubtedly powerful, it relies on a certain amount of experience and comprehension of general Linux terms.

Kudzu, started at boot-up, checks for new hardware and tries to configure it. Also, the new *RP3* utility for creating a dialup Internet connection is accomplished and simplifies a complicated process. These, together will all the other small utilities like *printtool* and *sndconfig*, make it a very manageable system that's fairly quick to set up.

However, while we didn't experience any difficulties ourselves, reports from other users describe problems with the installer freezing or generating errors, and there have apparently been complications when upgrading. In general though, Red Hat Linux is a powerful and flexible distribution which has excellent documentation and a very useful collection of configuration programs.
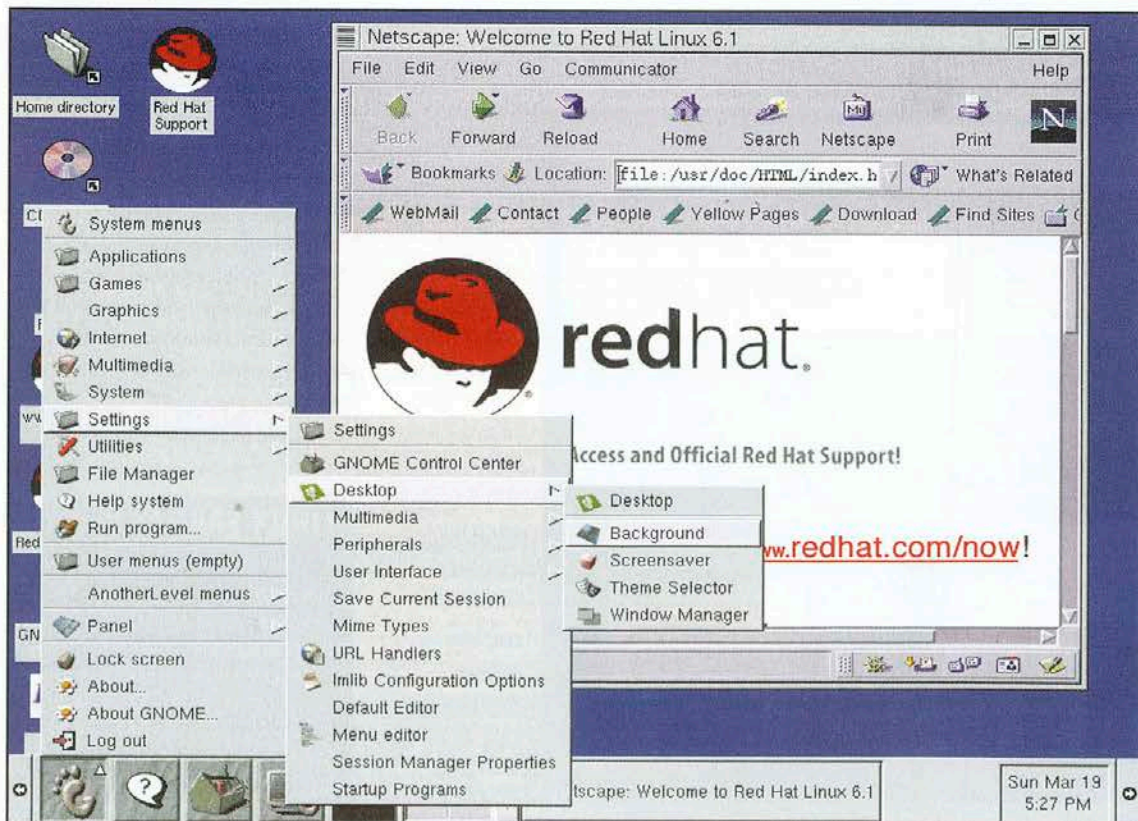
LINUX FORMAT Verdict

Ease of use:	8/10
Documentation:	10/10
Features:	9/10
Software range:	8/10

Robust, with a solid range of software. However, there are a few installation problems.

Rating 8/10

Debian

It may be a bit behind the times, but Debian is one of the best distributions on the market.

Debian GNU/Linux 2.1 – http://www.debian.org

dselect, *Debian's package management system, listing the installed software.*

Unlike the other commercial distributions in this review, Debian is produced by a non-profit making organisation called Software in the Public Interest. Volunteers all over the world help to maintain the system, and the distribution has acquired considerable respect from experienced users in the community. The text-mode installer is available in mono or colour, so Debian can be installed on very old systems that don't have colour displays. This is a good idea, and a refreshing change from the new, jazzy installers we're seeing everywhere else.

Are you experienced?

From the very start, it's clear that new users would find this heavy-going. The *cfdisk* partitioner assumes plenty of knowledge and there's little of the hand-holding you get with others. Conversely, you're constantly prompted and alerted as to what the installer is doing, which is a world apart from the 'do everything automatically' idea behind distributions like Corel.

The installation takes some time as after choosing one of the package categories (desktop, development workstation, etc), you have to sit through the post-install configuration and occasionally answer questions for each package. Debian uses the .deb format, which is similar to RPM. However, with the extensive *dselect* package management tool, adding and removing software is straightforward.

The best feature of all is the *apt* (*Advance Package Tool*) system, which allows the entire system to be updated with just a few commands. After experiencing incompatibilities with RPMs from various distributions, we
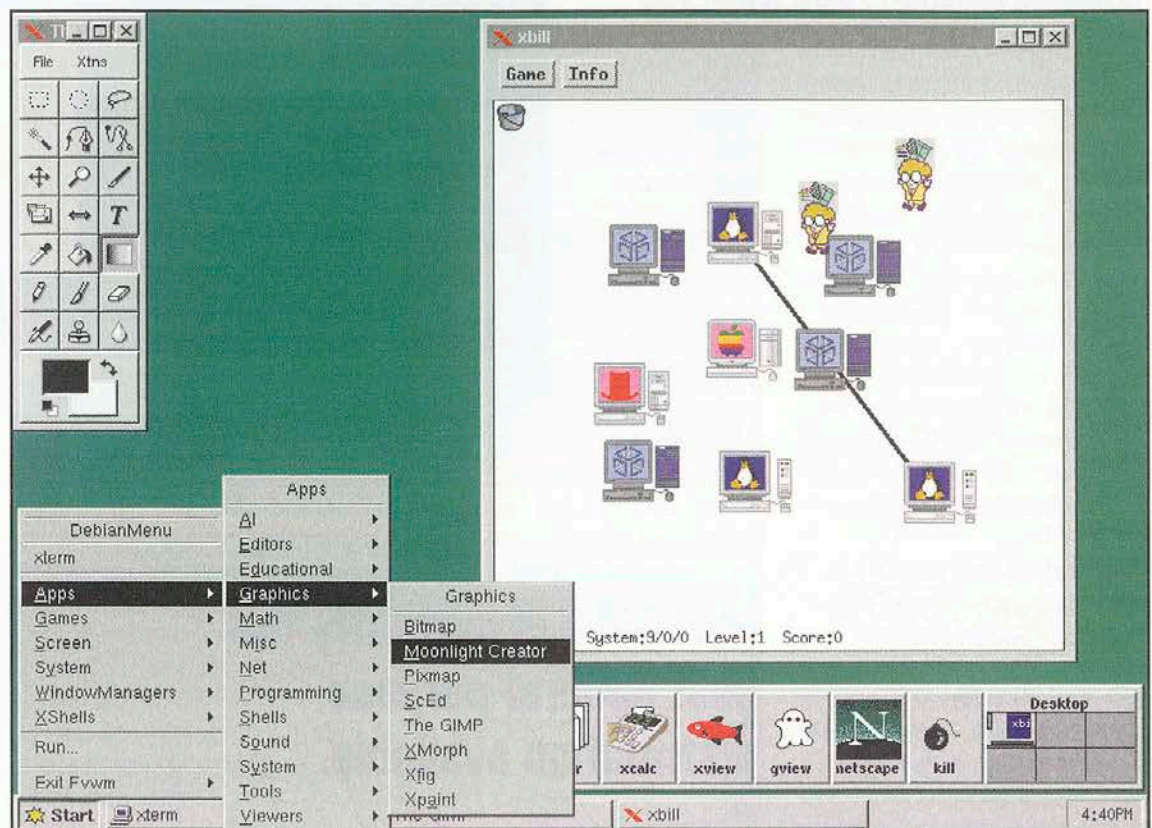
were glad to see such painless upgrading and package control.

On the software side, Debian is showing its age. The supplied kernel is 2.0.38 – not even from the 2.2.x series – and *XFree86* is back at 3.3.2. An old version of *GNOME* is supplied, along with several window managers, such as *Enlightenment, Window Maker* and *FVWM*.

Being a technically-orientated distribution, many console programs that don't need *X* are included, with mail programs, development tools and network utilities.

The Debian CD's packages are slightly behind the times for a few

reasons. Firstly, as it's a volunteer project, the developers can't just throw out another release whenever possible. More importantly, they are concerned with the stability and maintenance of a working system, as seen by the excellent upgrade process.

This development process and general attitude is similar to that of the kernel itself – technical excellence takes preference over commercial interests. Debian is not a good choice for Linux beginners, and if you're starting out for the first time, you'll have an easier introduction with Mandrake or Caldera. However, for experienced Linux users who need a

FVWM 95 is a fast and small window manager, emulating the look of Windows 9x.
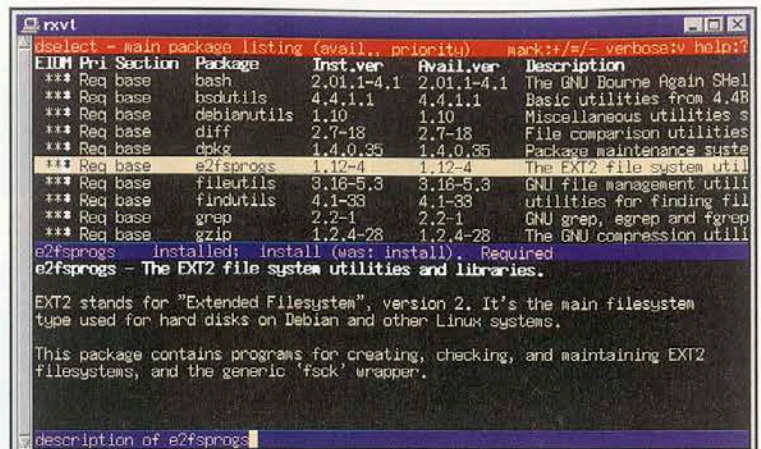
very reliable, easily maintainable and carefully crafted distribution, this is the best you can get.

LINUX FORMAT Verdict

Ease of use:	4/10
Documentation:	9/10
Features:	9/10
Software range:	6/10

Not suitable for new users, but technically superb and easily maintainable.

Rating 9/10

Linux-Mandrake

Simplicity and a wide range of applications make this a great distribution for all Linux users.

Linux-Mandrake 7.0 – http://www.linux-mandrake.comhttp://www.redhat.com

Mandrake's desktop, KDE, as it appears at first boot.

The concept behind Linux-Mandrake was originally born out of Red Hat's refusal to include *KDE* in their distribution. This was due to licensing concerns, so Mandrakesoft decided to create a variant of Red Hat with *KDE* included and ready to go. This meant new users wouldn't have to find and download the desktop suite from the web. Since then, Mandrake has evolved into a solid distribution in its own right. The new graphical installer is a fantastic, easy to use tool, and we had absolutely no problems getting underway with it.

A lot of thought has clearly been put in here: a column of lights down the side acts as a progress meter, while a bar of help text at the bottom provides tips for choices in each dialogue box. You can choose a preset security level and even decide if num-lock is activated at boot time! Thankfully, a proper user account is created and PPP (dialup) configuration is done here too. Apart from the long package installation process and lack of ability to configure sound, Mandrake's installer is the best graphical one we've seen and it's a joy for both newcomers and more experienced users alike.

In operation, the distribution still remains similar to Red Hat. However, one major enhancement lies in the hardware setup tool, *Lothar*. This automatically detected most of our hardware, although configuring it still proved to be difficult. Still, it's certainly an impressive start and with more development it could really attract non-technical users to Linux. The *DrakConf* configuration tool will also be a great help for newcomers and assists in a range of administration chores, while the *Supermount* feature, for automatic mounting of removable media, is a nice touch.

On the desktop side, Mandrake has *KDE 1.1.2* by default, and *GNOME, Window Maker, IceWM* and *Enlightenment* are included as well. Also on the CD are *Netscape*, the *GIMP* and *WINE* for running Windows apps, while *Emacs, KDevelop* and *egcs* are available for developers.

Mandrake's *KDE* setup is good, and the on-disc documentation is thorough and well-written. The range of apps on the CD is varied and up-to-date, and while it's doubtful how much of an effect the Pentium-class processor optimisations have, the system feels snappy and responsive.
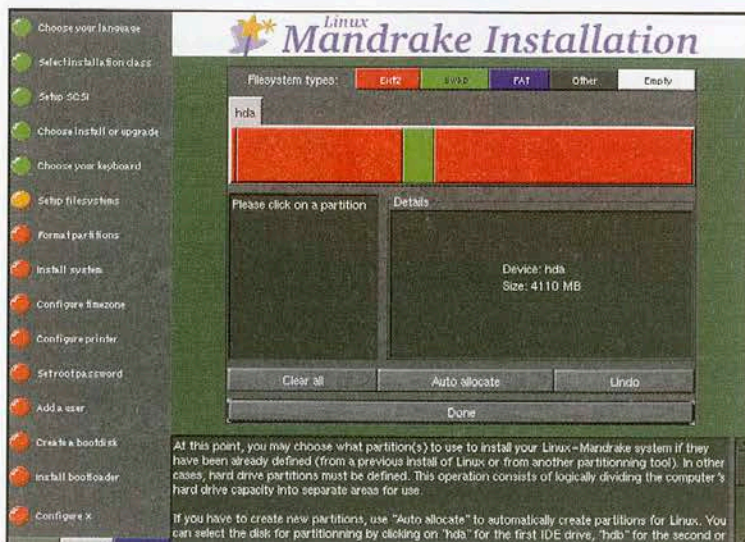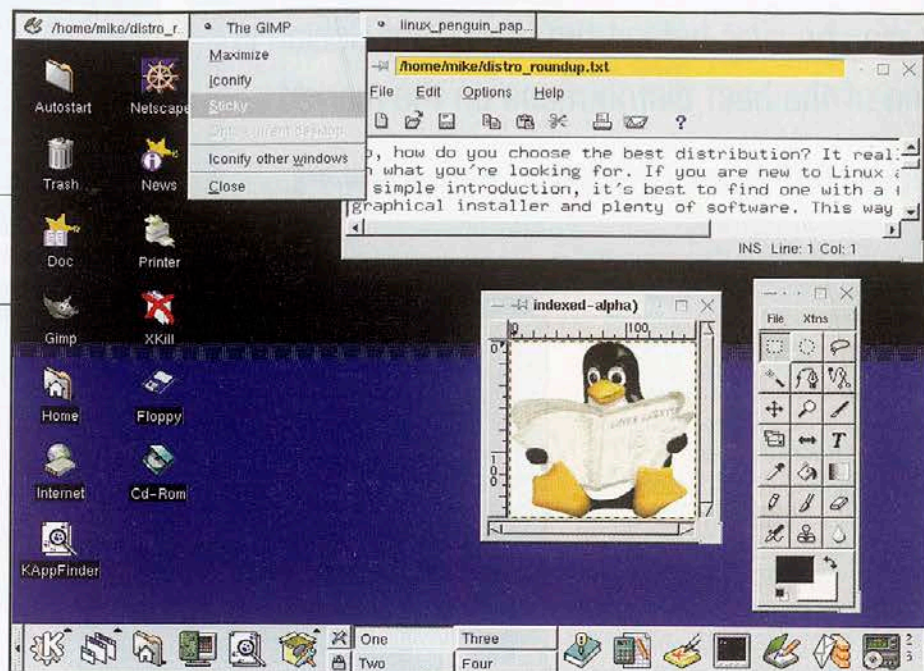
The greatest?

With all this in mind, and the brilliance of the *DrakX* installer, Mandrake is currently the best general-purpose distribution available. New users will love the easy setup process, ready-to-go *KDE* desktop and helpful documentation, while experienced Linuxers will be pleased with the variety of desktops, development tools and general flexibility of the system.

Mandrake is currently the best general-purpose distribution available.

The excellent installer, with its powerful partitioning tool.

LINUX FORMAT Verdict

Ease of use:	10/10
Documentation:	8/10
Features:	9/10
Software range:	10/10

Easy to use, robust, fast and with lots of great software. Highly recommended.

Rating 10/10

Caldera OpenLinux

If you're in the commercial business sector, this may be the perfect distribution for you...

Caldera OpenLinux 2.3 – http://www.calderasystems.com

Caldera's main focus with their OpenLinux distribution has been the commercial sector, providing a stable desktop and server operating system that fits well into a business environment. It can be installed through Windows or onto a clean hard drive and sports a fully graphical setup process. *Lizard*, the installer, is very polished and features a column of relevant help text down the side of the screen, which is invaluable for newcomers and is a useful reference as the installation progresses.

Our mouse and graphics card were detected correctly, and the partitioning tool was clear and straightforward to work with. We could create a default user and finish some other configurations while the packages were being installed, which is a nice touch, and saves wasting time too. However, we came across a snag when starting for the first time because the swap partition hadn't been assigned properly. Although we could change the fstab file and add a swap entry ourselves,

COAS, Caldera's administration tool, in action.

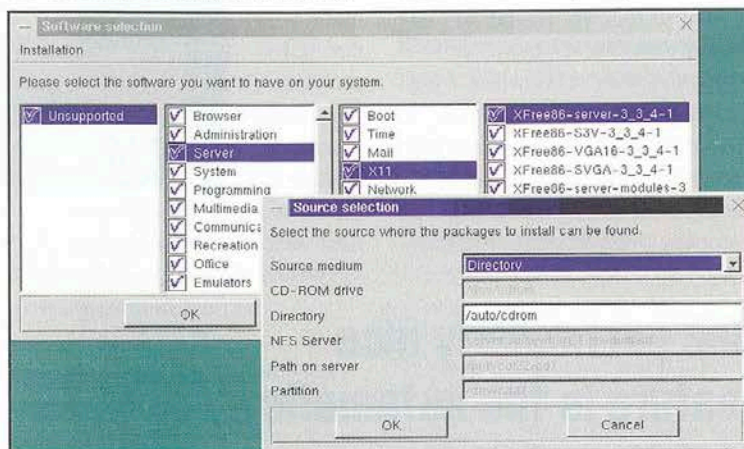new users would find the system very sluggish and confusing.

OpenLinux's default desktop is *KDE* which, when started for the first time, launches *Kandalf*, a wizard that offers various customisation options for the desktop. Overall, it has a polished and professional feel to it.
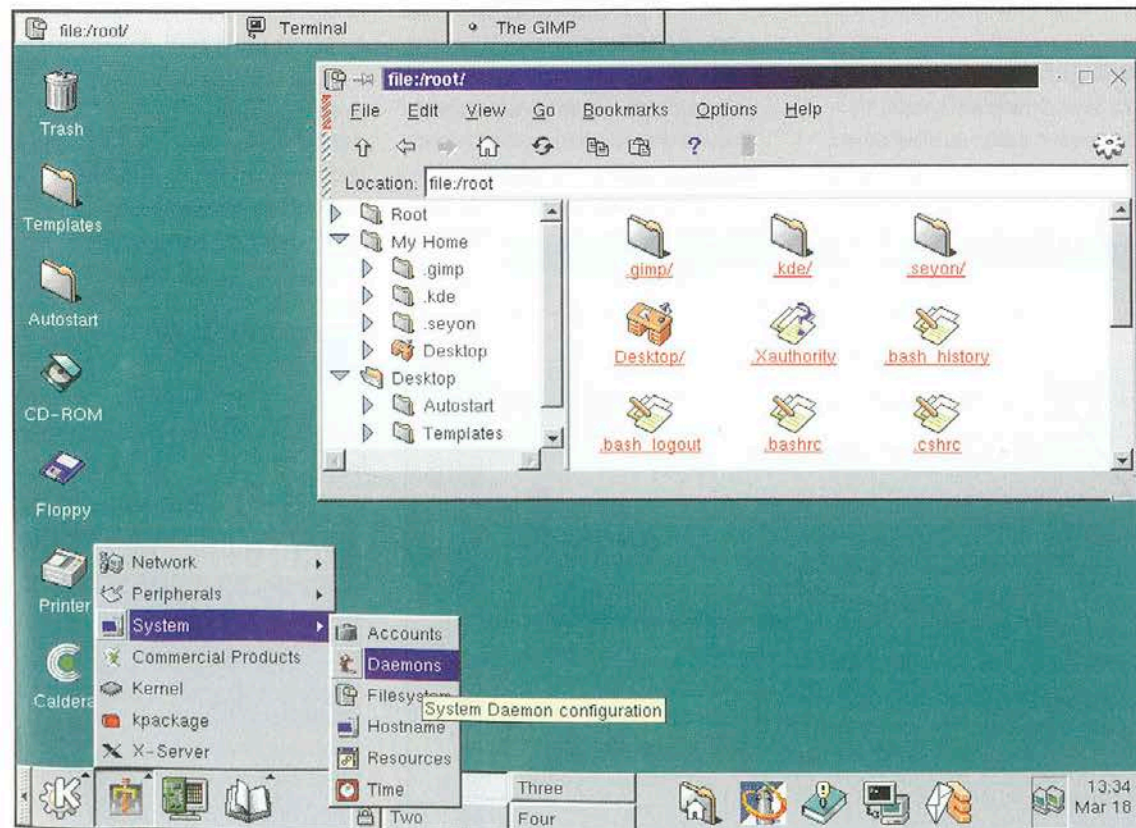
COAS

COAS, or the Caldera Open Administration System, is a suite of programs designed to make configuration and management of the system easy. While not as comprehensive as *linuxconf*, it's less daunting and simpler to work with. COAS includes utilities for managing users, setting up a network and installing packages, among others.

The usual desktop applications like the *GIMP* and *Netscape* are installed by default, while *Dosemu*, *WINE* and *Samba* are provided for Microsoft OS connectivity. Developers are catered for too, with *XEmacs* and *egcs*.

Apart from the swap partition difficulty, we found OpenLinux to be a very competent distribution, suited for newcomers and intermediate users alike. There are some problems with using RPMs and most veteran Linux users won't feel at home here, but overall it's a polished distribution and it makes a good desktop OS all-round.

The plain KDE desktop, as set up by the Kandalf configurator.

LINUX FORMAT Verdict

Ease of use:	9/10
Documentation:	8/10
Features:	9/10
Software range:	8/10

Polished and user-friendly, OpenLinux is a good choice for the desktop.

Rating 8/10

SuSE Linux

We test this long-established German company's latest distribution effort. Can it hold off the competition?

SuSE Linux 6.3 – http://www.suse.com

As experienced distributors, German company SuSE have been around for many years and have been very popular in the European markets. One of the reasons for this is undoubtedly due to the large amount of software found on their CDs. They have also recently converted their distribution to run on the Alpha processor, and they have put money into several of the key areas of Linux development.

SuSE's installer is based on *YaST*, a long-running setup utility which has finally been updated from text-mode to a full graphical installation. Sadly, when we started our *YaST2* installation, it told us that our 2GB partition was too small for the minimum installation. Wisely, though,

The YaST tool assists with the administration and configuration of your hardware.

SuSE have included the original tried-and-tested *YaST*, and this worked properly. Installation is reasonably quick, although it isn't very neat – first *YaST* starts, then exits to restart the system services, then starts again, then exits and waits for a minute before running some scripts. However, it's a reasonably simple installation and the online help is thorough and useful.

Configuration tools

YaST is also the main system configuration tool in SuSE, offering different administrative and setup options, from adding new users to

configuring modems. For the most part it's easy to use and powerful, although occasionally it didn't do exactly as commanded, and if you're editing configuration files by hand it can have trouble working with them. Like others, SuSE has opted for *KDE* by default. It is set up fairly well, with icons for the disk drives and easy access to *YaST* and the help system. One particularly nice feature is the automatic updating of menus – install a new application through *YaST* and it will add an entry to the *KDE* menu. This can also work for the *FVWM* and

One particularly nice feature is the automatic updating of menus...
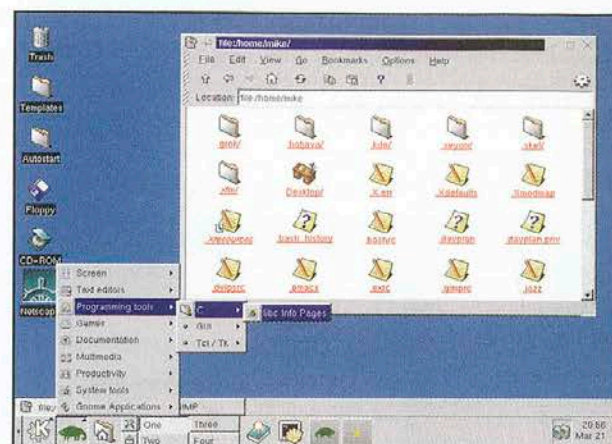
KDE with the SuSE menu, automatically generated from the installed applications.

IceWM menus (among others) as well. SuSE Linux uses the *RPM* package format, but we had compatibility troubles installing many off the web. Thankfully, SuSE provide an excellent service on their site for new *KDE* programs, which are tailor made to fit the distribution.

One cause for concern we noticed was the inclusion of *GIMP 1.1.11*, a development version. The program's docs say that it's an unstable version "intended for developers only", so it seems SuSE are chasing version numbers here. This is a shame, but a good, wide range of software for both the desktop and server is included, which makes up for it.

In all, SuSE is a generally good distribution with some attractive features. Unfortunately, some of its quirks hold it back slightly.


```
 – YaST
                    YaST - Yet another Setup Tool
                YaST Version 1.03 -- (c) 1994-99 SuSE GmbH

    Language:     English
    Media:        CD-ROM ATAPI EIDE /dev/hdc
    Root-Device:  /dev/hda1

    General help fo  Integrate hardware into system    ->
    Adjustments of   Kernel and bootconfiguration      ->
    Choose/Install   Network
    Update system    Configur  Network base configuration
    System administ  Login co  Change host name
    Show README fil  Settings  Configure network services
    Copyright        User adm  Configuration nameserver
    Exit YaST        Group ad  Configure sendmail
                     Create b  Configure ISDN parameters
                     Security  Configure a PPP network
                     Set the   Administer remote printers
                     Set time
                     Configur
                     Configur
                     Change c
```

LINUX FORMAT Verdict

Ease of use:	8/10
Documentation:	8/10
Features:	8/10
Software range:	9/10

A decent distribution with some clever features and a wide choice of software.

Rating 8/10
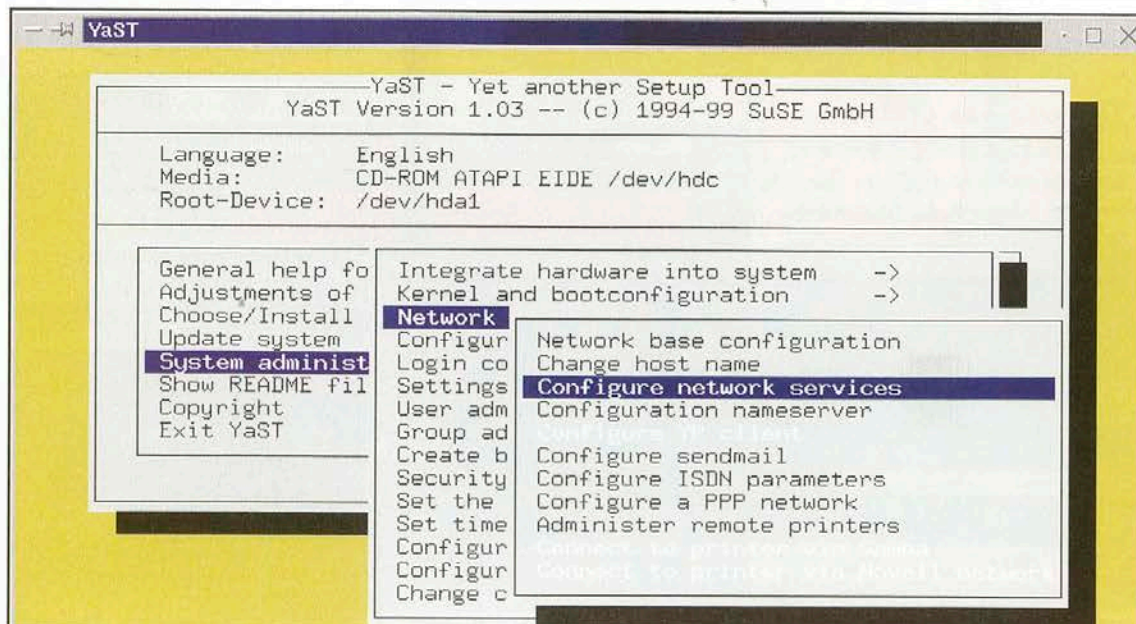
Conclusion

Of the eight distributions on test here, some are clearly more suited to different types of users than others. Corel, for instance, definitely have newcomers to Linux in mind with their attempt at automatic hardware detection and simple install routine.

As most unfamiliar users will be coming from a Windows background, extra effort has been made to provide a familiar working environment with an *Explorer*-like file manager and integrated video and network configuration in the desktop control panel. On the other hand, distributions like Debian and Definite are aimed at the more experienced users who will be putting reliability and flexibility above flashy desktops and installers.

This is important when making your choice – you'll want an easy introduction if you're unfamiliar with Linux, but if you know your way around then the restrictions of novice distributions may be frustrating.

Installation issues

While most are starting to feature graphical installers (based on the *VGA16 X* server for compatibility), this isn't always the best choice. Linux is proving to be increasingly popular on older hardware, where it can function as a file or print server, or simply as a development OS for trying out new software. In this light, text-based installers are crucial – running a graphical installer would be far too slow or even impossible. The best solution is to provide both, as Red Hat have done, so you can revert to the standard setup program if problems occur. In general, the graphical installers in these distributions add little to the functionality, but provide an important level of comfort for novices.

Another issue to watch out for is with package formats. Most distributors have opted for the RPM format created by Red Hat, but this doesn't guarantee compatibility with others. RPMs made for some systems may not work on others, as we've seen with OpenLinux and SuSE. If you're careful, you can avoid checking for dependent packages when installing (using the '-- nodeps' switch at the command line), but this may just cause more problems.

In general, the distributions derived from Red Hat – such as Mandrake and Definite – have less trouble with using standard RPMs off the web. SuSE and Caldera aren't left out, but finding packages that will work with these isn't as simple. Their sites can provide tailor-made RPMs that are ready to use, but they're not always totally up to date.

Corel and Debian take a different approach by opting for .deb packages, although you can use the *Alien* program to convert RPMs to the default
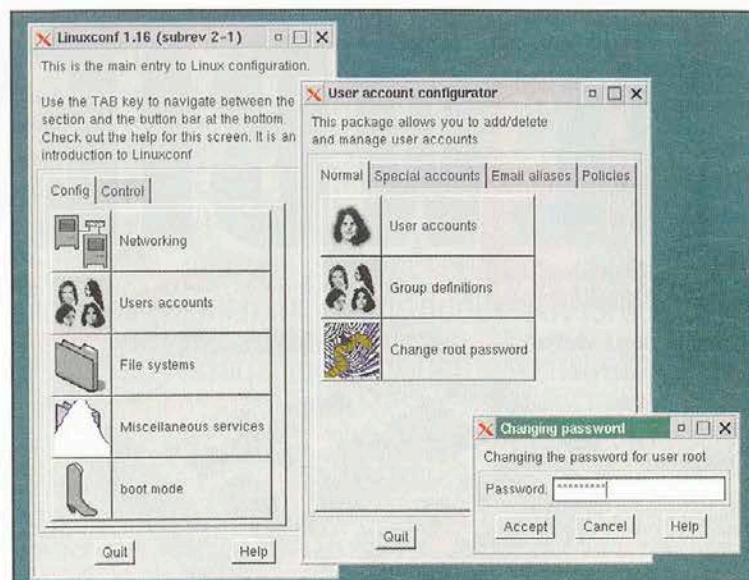
format. While RPM packages are more abundant around the net, there's no shortage of .deb archives, thanks to the popularity of Debian.

Desktop decisions

Pop in to a Linux newsgroup or join a mailing list and you'll see heated debates over *KDE* and *GNOME*. While *KDE* was started earlier and was the first to have a solid, usable desktop environment, *GNOME* has caught up and now there's not a lot between them. Because they're both freely available, you can try them out yourself and see which one you prefer.

If you don't like either, or want to run something less resource-hungry, you can switch to a simple window manager like *FVWM* or *Blackbox* and use most of the *KDE* and *GNOME* applications from there.

Still, it has to be said that most distributions reviewed here feature *KDE*

linuxconf is the standard setup tool with distributions based on Red Hat.

as standard – it's slightly easier to work with in places, and due to the native window manager (*kwm*), new users don't have to worry about two different configurations at the same time. However, with a compliant window manager, *GNOME* provides an equally friendly drag-and-drop desktop with lots of useful utilities.

And the winner is...

From our tests, we found Mandrake 7 to be the best of the bunch out of the eight reviewed here. Debian and Definite are perhaps better alternatives for power users and experienced Linux fans, while SuSE, Caldera and Red Hat are all worth considering as well.

However, if you're looking for a friendly, powerful all-rounder with lots of great software, Mandrake really should be your first choice. **LXF**

TABLE OF FEATURES

	Kernel version	XFree86 version	Auto X setup	Glibc version	Main package format	Graphical installer	Default desktop	WINE to run Windows apps	Compiler	Boxed-set support
Corel	2.2.12	3.3.5	yes	2.0.7	Deb	Yes	Modified KDE	No	gcc 2.7.2	30 days
Definite	2.2.12	3.3.5	Yes (Xconfigurator)	2.1.1	RPM	No	GNOME 1.0.10 pre2	No	egcs 1.1.2	60 days
WinLinux	2.2.13	3.3.5	Yes (in Windows)	2.1.2	TGZ/RPM	Yes (in Windows)	KDE 1.1.2	No	egcs 2.91.66	Installation only
Red Hat	2.2.12	3.3.5	Yes (Xconfigurator)	2.1.2	RPM	Yes	GNOME 1.0.39	Yes	egcs 1.1.2	90 days
Debian	2.0.38	3.3.2	No	2.0.7	Deb	No	FVWM	Yes	gcc 2.7.2	N/A
Mandrake	2.2.14	3.3.6	Yes	2.1.2	RPM	Yes	KDE 1.1.2	Yes	egcs 1.1.2	100 days
[Caldera OpenLinux]	2.2.10	3.3.4	Yes	2.1.1	PRPM	Yes	KDE 1.1.1	Yes	egcs 2.91.66	90 days
SuSE Linux	2.2.13	3.3.5	Yes (SaX)	2.1.2	RPM	Yes	KDE 1.1.2	Yes	gcc 2.7.2	60 days

NEWBIES START

Don't be nervous about installing Linux, it's easier than you imagine. In combination with the version of Definite Linux on our CD, our special beginners' guide will soon have you running Linux!

New to Linux? Maybe you already know a bit about Linux, but are unsure about how to install it, or perhaps you're just curious as to what all the fuss is about. Either way our guide will have you up and running in no time. Well, maybe an hour or so...

WHAT IS LINUX?

Linux is currently the fastest-growing operating system in the world (in number of users – not size!). Why so much fuss about an operating system? Can it really be so fantastic? Well, don't take our word for it – try it out for yourself. First, here's some background information.

Linux is an operating system – that's the collection of software on your computer that actually detects and controls all the hardware and provides a user interface for you to run your programs. Other examples of operating systems are Windows 98, Windows NT, MacOS, BeOS, AmigaOS, FreeBSD and Unix.

You probably haven't heard of some of these, but the most important one on the list is probably Unix. Unix is a powerful operating system with quite a history, and it's normally found driving huge mainframes at universities, or large industry servers. Linus Torvalds came up with the idea of Linux because he wanted to run Unix-style applications on his home PC. He couldn't afford to buy a version of Unix and it wouldn't run on his PC anyway, so he set about creating his own version. He also decided to make it an "open-source" project, which meant that the results of his work would be available to everybody.

WHY IS IT IMPORTANT?

There are several reasons why Linux is important. Firstly, it is "free" and open source. This means that it doesn't take a lot of money to set up a powerful Linux system, and all the upgrades are freely available. Linux is in a constant state of development, and whenever you want to update the main parts of your system's software, you can just download them free!

Another advantage of Linux is, because of its open source nature, it is a very good platform for developing software on. There are countless programming tools available, and countless programmers to help you when you get stuck. You'd probably be

Linux Info

"Newbie" is a friendly term which just means someone who is a beginner at something. There is no stigma attached to being a newbie – Linux users are a friendly bunch on the whole, and they will be quite happy to help you out and give advice if they know you're a newbie.
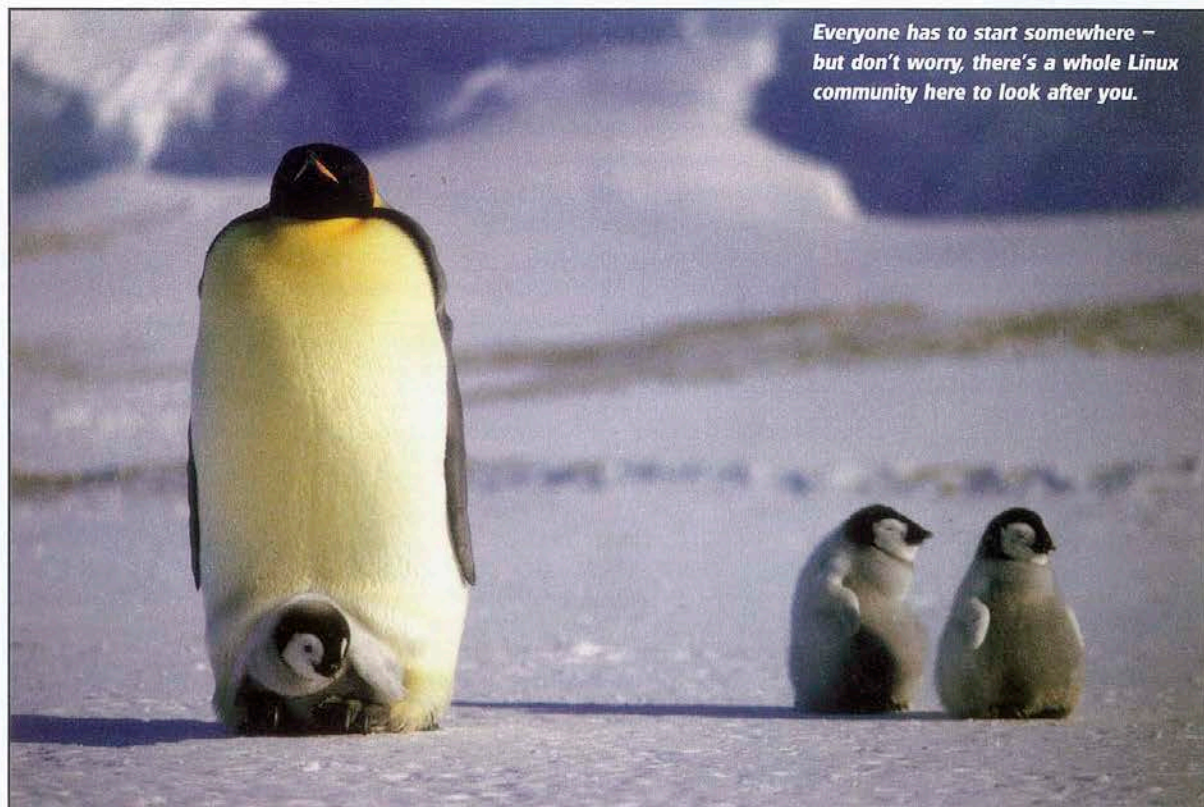
Everyone has to start somewhere – but don't worry, there's a whole Linux community here to look after you.

HERE!

surprised to discover that, for example, the game *Quake* was developed on a Linux system.

Perhaps the most important thing for some people, especially those who can't afford to upgrade their hardware every six months, is that Linux is quite lean in terms of the resources it requires. You can quite happily run Linux on a 486 and you aren't locked into an expensive upgrade cycle.

The best thing of all is because the entire system is open-source, you can tinker with any bit of it you like. You might need to know a bit about programming before you start writing device drivers, but every bit of the operating system and most of the software that runs under it can be tinkered with to make it do exactly what you want it to. Hurrah!

The downside to Linux is that it is perhaps not as transparently easy to use as, for example, a Mac. At *Linux Format* we think it is a good thing for computer users to understand a little about how their computer actually works, but we understand that this is perhaps not for everyone. Things are changing though, and much work has been done on creating GUIs (Graphical User Interfaces) to make Linux easier to get to grips with. The most popular of these are *KDE* and *Gnome*, and they are supplied with almost all distributions. So what's a distribution?

DISTRIBUTIONS

Because of the open-source nature of Linux, people have been positively encouraged to create their own unique versions of the operating system, adding different methods of installing and maintaining it, and building in unique features. The result is that there are dozens of different distributions (or "distros" for short) available. The most popular ones are Red Hat and Mandrake at the moment, but Corel have also created a new distribution which they claim is the easiest to use.

For more on the different distributions, check out the roundup in the Reviews section of this issue.

You might think that all these distributions lead to a dangerous fragmentation of the market, where everything eventually falls apart and becomes incompatible. This doesn't happen for a number of reasons:

• The kernel, the main part of the operating system, is essentially the same in all distributions.

• Because almost all Linux software is open-source, it can be distributed as source code and compiled on the user's own computer, ensuring compatibility.

• Most of the distributions are closely related. For example, Definite Linux and Mandrake are very similar to Red Hat.

WHAT IS THE KERNEL?

The most important part of Linux is the kernel. This handles the computer's most fundamental input/output operations, schedules processing, deals with the memory and contains the major hardware drivers apart from graphics.

The kernel is the most frequently upgraded part of Linux and the same kernel is common to all the different flavours of Linux around. It is able to maintain this coherence because its definitive form is edited and tweaked by Linus Torvalds himself from submissions from a number of kernel developers around the world.

New versions of the kernel are periodically released over the Internet and so it's important to know which kernel you have and how to upgrade it when a new one appears. Currently, the most up to date stable version of the kernel is 2.2.14, but version 2.4 is due out really soon, and it's probably already being downloaded as you read this feature. The latest version is always available from **http://www.kernel.org** or one of its many mirror sites.
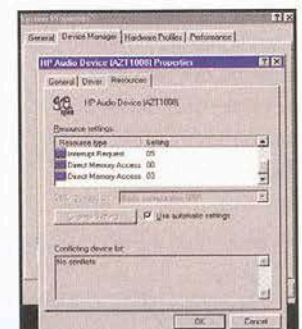
To begin with, however, it isn't particularly essential to have the very latest version of the kernel, and you'll no doubt have enough on your plate to begin with, without having to worry about installing a new kernel.

People have been positively encouraged to create their own unique versions of the OS

Who said Windows was good for nothing? The System control panel holds some useful info.
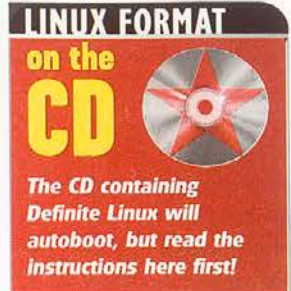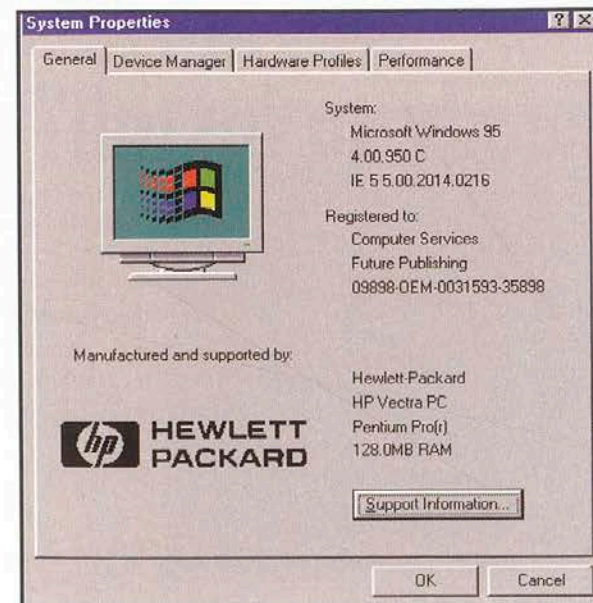
IRQ numbers for sound cards and other useful info should be jotted down.

Pretty exact descriptions of your hardware can be found here.

Ready to begin?

Now you know a little bit about Linux, it's time to get ready to install it. You can just ignore these instructions if you like, boot up the CD, randomly answer questions and get on with it. As with installing any operating system though, there is a distinct possibility that you will completely mess up your hard drive and lose any data you might have, so perhaps you'd better read our instructions, which will give you a much better chance of success. We don't want to scare you because it shouldn't be a problem, but the better prepared you are, the more likely it is that you'll avoid losing your data.

For the purposes of this guide, we will assume that you are installing Linux in addition to another operating system (Windows, most likely), and that you will want a 'dual-boot' system – one which will ask you which OS you want to use when you turn it on.

Now you obviously know better than we do what hardware your system is running on. Even if you know pretty much what all your bits are, it is a very good idea to make certain, just so that you are able to answer any questions you may be asked during the install.

An easy way to do this is to boot up into Windows – it may be bloated and clumsy, but it does know a lot about your hardware. Open up the control panel and double click on the System icon. This lists all your hardware that Windows has found. If you click on an item and select properties, you'll find out a lot of extra information.

A lot of the hardware will be detected automatically, but it's always handy to have the specifics in case something goes wrong. Things that you will need information on are your:

Monitor: The best information for this will come from the original documentation, or the manufacturer's website. Basically, all you'll really need to know is what refresh rates and resolutions it can handle.

Graphics card: Find out as much as possible about your card. The type will be enough if it is a common one. You should also note down how much memory the card has, what type of chip it uses and other details you think may be important.

Memory: Knowing how much memory you have is helpful for setting up swap space, etc.

Network card: If you are connected to a network, you should note down what network card you have. You should also make a note of your IP address, gateway and DNS servers from the Network control panel, and the netmask value. If you are setting up Linux on a LAN, you can get this information from the network administrator too (if that isn't you), and it might be nice to let him/her know what you are up to!

SCSI interface: If you have one, make a note of the type. Pretty much every 'proper' SCSI card is supported, but you may have problems with ones that have been "given away" with cheap scanners. These are usually cut down generic cards which don't support proper SCSI commands.

Printer: The model of the printer, and the port it is connected to, would be handy.

Disk Drives: These will, for the most part, be auto-detected by Linux, but it will help you considerably when you are partitioning the drives to know which is which and where they are located (for example, for IDE drives, which are slaves and masters).

Modem: If you have an external modem, it is likely to be supported. To help you get the best use out of it, make sure you know what type it is (V90, K56flex, etc) and if it is based on a standard type of modem hardware (for example, Rockwell).

Once you have a list of hardware, it may be an idea to check that it is supported. You can find out on the Red Hat website: **www.redhat.com/support/hardware/** Everything listed under RedHat 6.0 will be supported by Definite Linux. Don't worry too much if your hardware isn't specifically listed because a lot of hardware devices use standard chips, so your device may be supported by another driver. For example, the Adaptec AHA2930 uses the same driver as nearly all of the other AHA SCSI cards.

This is particularly true of graphics cards, since all modern cards will work in a straight SVGA mode, although they will not be as fast as they would using an accelerated driver.

Ready to install?

Now you are ready to install – or are you? There are two more things we need before you start: somewhere to install Linux and a way to boot up the installer.

The first step is usually the one which causes most concern. There are two possibilities here – that you are installing onto an unused hard disk (or one with unused partitions), or you have a disk full of Windows partitions and you want to squeeze them to make space for Linux.

The first option involves less risk to your Windows data (ie, virtually none). It can be a good idea to have Linux on a separate disk as this makes it easy to re-install or remove. If you have, you can skip ahead to Partitioning Your Drive. Failing that you will have to resize your Windows partitions. First, check how much free space you have on your Windows partition. You might want to clear out some junk and work out how much space you want to leave Windows with. For average desktop use, with the usual aps and window managers, with space for files, etc, you are going to easily need 600Mb of space for Linux. We would recommend 2Gb. If you have nowhere near that sort of space, now might be a good time to buy a bigger drive.

Now you should run *Scandisk* on your drive to fix any errors and defragment it. Make sure you choose a Full Defragmentation option because you want all the files to be relocated to the front of the drive. (If you are using Windows98, uncheck the option 'Rearrange program files so my programs start faster' to achieve the same effect). This process can take from a few minutes to several hours.

Once you have performed all these tasks, you'll be ready to resize partitions. Windows doesn't come with any software for this task (surprise surprise), but there are free tools on our ➔

USING FIPS TO RESIZE YOUR DRIVE

The aim of using FIPS is to create space that will then be used for the Linux partition. Thousands of people have used this system without any problems to set up Linux, but this is the part of the process most prone to data-threatening problems, so make sure you are familiar with the procedure before you start. FIPS is provided without any warranty and Linux Format cannot give support for it, so read the instructions!

Step 1

The most important file is the FIPS.doc. Read it. Print it out. Read it again. Make sure you understand the process before you proceed.

Step 2

If you haven't backed up your crucial data, DO IT NOW. This is also a good time to create a DOS or Windows rescue disk, just in case. (You should have one of these anyway.)

Step 3

Make another bootable floppy. FIPS will not be able to work on your Windows partition if you are running under Windows. Now copy the files FIPS.EXE, RESTORRB.EXE and ERRORS.TXT to the boot floppy. You might also want to copy the doc files if you haven't printed them out – if an error occurs you might not know what to do without the docs. When your data is backed up and you are ready to go, restart your computer and boot from the new floppy.

Step 4

When you arrive at the A:\>; prompt, type FIPS and press Enter. The first thing you will see is a warning about using FIPS in multitasking environments such as Windows. As we booted from a floppy we are safe here, so press Enter. Next, FIPS will analyse your existing partitions. It will probably pause for a long time at the 'Checking FAT' and 'Searching for Free Space' stages. This is quite usual because we are dealing with large amounts of space here.

Step5

When FIPS is done with its analysis, it displays the results. At this point you may get a warning of something wrong with your FAT. The most likely cause is that you have a large hard drive. Read the message carefully and you will find that this is normal with large hard drives and won't prevent FIPS from working properly. Remember to always read the messages carefully.

Step 6

FIPS will now show you how it plans to split the existing partition and you will have the opportunity to make changes. By default, FIPS appears to take two thirds of the available free space for the new partition it creates. You may want to adjust the partitions to allow Windows to keep more space, especially if you plan to install any new software on Windows. Adjust yours based on your own needs, but we'd recommend using at least 1Gb for each operating system, space permitting. Use the Up and Down arrow keys to make large changes (10 cylinders at a time) and the Left and Right arrow keys for small ones (one cylinder at a time). The size of the existing partition is shown on the left and the size of your new partition is on the right. In the middle is the cylinder number where the split will take place.

Step 7

When you are satisfied with your new partition plan, press Enter. FIPS displays information about the new partitions and asks permission to write it to disk. Your hard drive has not been altered at this point. You may choose to write this configuration to disk or re-edit the partition table, which starts the process over from the beginning. If you get an error message about missing files when you try to re-edit the data, don't worry. Just reset with CTRL-ALT-Delete and start over – you won't have damaged your drive.

Step 8

When you write the new partitions, FIPS will offer to make a backup of your existing boot sector and we recommend that you do this. The backup file is tiny, made in a second and will be invaluable if anything goes wrong.

Step 9

After FIPS completes its work, it displays another message instructing you to run Scandisk on your old partition (Windows sometimes miscalculates the used and free space on your drive after using FIPS, but Scandisk corrects this problem). If you choose to restore your original partition scheme using the RESTORRB utility, you should run Scandisk after this as well. Sometimes you may receive an error like this: 'Memory Allocation Error, Unable to Load COMMAND.COM'. Don't be alarmed if you see this – just press Control-Alt-Delete to reboot and all should be well.

Step 10

FIPS has made space and has created a second partition using this space. Both partitions will be designated as 'Primary'. The second partition that we've just created is now taking up the space on the drive where you want to install Linux. When you run the Linux installer, it will want to create its own partitions in this space, so it is easier if we delete this partition now (the space is the thing we're interested in, not the partition itself). Do this now by booting from a DOS boot floppy and running the DOS utility FDISK from the DOS prompt. Note: this is not the same utility as the similarly-named Linux utility fdisk. How the partitions will be displayed depends on your current configuration, but you should be able to determine which is the partition created by FIPS, as it will be after your Windows partition. Be careful, though! Do not delete the first partition on the drive as that is where Windows lives. If you do, you will lose all your Windows programs and data.

Step 11

Just to reassure yourself that everything has worked, you'll probably want to reboot from your hard drive and run Windows. This is a good time to run Scandisk on the drive, so just to make sure it works. Now you are ready to install a proper Operating System!
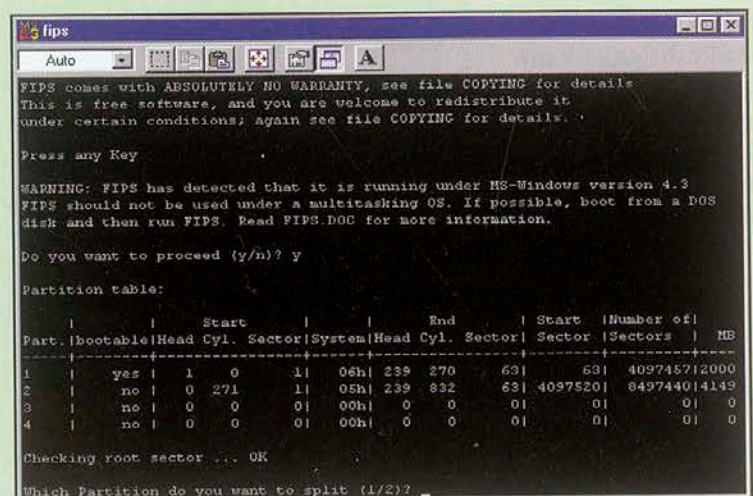
→ cover CD to help you. Check in the dosutils drawer on the disc, and you'll find a DOS utility called *FIPS* (*First Interactive Partition Splitter*).

There are several commercial products which are a good deal easier to use, such as *Partition Magic* and *System Commander*, which we have reviewed this issue. If you want to use *FIPS*, check out the step-by-step guide.

BOOTING UP WITH THE INSTALLER

The version of Linux on our CD will autoboot. However, if your PC will not boot from a CD (the easiest way to check is to stick our CD in the drive and reboot), then you will have to create an install floppy. This is pretty easy to do as there is a utility on the disk to do it for you.

The utility is *rawrite*, and you can run it straight from the CD (if your version of Windows will let you). First, format a blank floppy disk. Now double click on rawrite.exe in the dosutils folder on the CD. In the DOS window that opens, you will see the program request the name of the image file.

Type: **[Drive]:/images/boot.img**

[Drive] is the drive letter of your CD drive. You'll then be asked where to write the image. Enter 'A' and then press return. The bootable disk image will be written out to your floppy drive (this will take a couple of minutes). When the drive light stays off, your disk is ready and you can quit from *rawwrite* – it won't give you an exit message or anything.

Partitions

To successfully run Linux, you will only need to create two partitions – a root partition and a swap partition. However, there may be several good reasons for creating more if you ever want or need them.
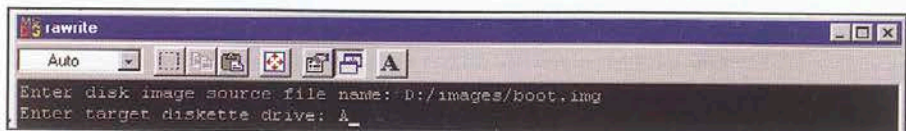
Rawrite *is pretty simple, but it works.*

One of the useful things about breaking up your Linux space into smaller chunks is that it makes it a lot easier to create backups, for example. Some of the useful ones to create are listed here.

THE SWAP PARTITION

The swap partition is a special type of disk partition which Linux uses to handle its virtual memory. You are probably familiar with virtual memory from Windows – basically, it uses storage space on the drive as memory. Virtual memory allows your computer to run large applications or work on files which are too large to fit into conventional RAM. It also means you can run more programs at once. The downside of this is that virtual memory is a lot slower than real memory, but Linux handles it cleverly so that the active tasks have access to the real memory, whereas the inactive ones are stored in virtual memory. As a result, the slowdown is limited to periods when tasks are "coming back to life", as it were.

Windows handles its virtual memory as a file on the Windows partition, which means that you don't have to set

RUNNING THE INSTALL PROGRAM

You should be ready to run the installer program now, so fire up your PC!

Step 1

Boot up your computer from the CD, or using the boot floppy if you created one. A message will appear on-screen, telling you that Linux is loading. Once the installer has finished initialising, you'll receive a welcome message and will be told of the options. Explore this information if you like, then press enter when you are ready to begin.

Step 2
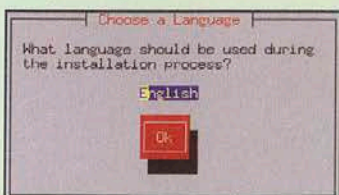
The installation will now begin. Nothing is actually being written to your disks yet, and it won't be until you specify partitions, so don't worry if you get confused and want to quit – just reset the computer. The first question is easy enough – it's a simple enquiry as to what type of keyboard you have. 'UK' is the correct answer for those who live in the UK, surprisingly enough.

Step 3
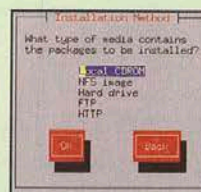
The next question is pretty easy too: where are you installing from? The answer is the default, Local CD-ROM. In the unlikely event that your CD drive isn't supported by Linux, you may have copied all the files on the CD to a hard drive, in which case you could simply select the HD option.

Step 4
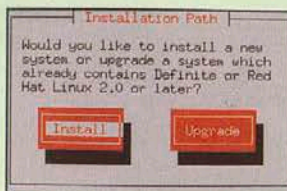
You will now be asked to insert the CD (which is most probably already inserted in your PC if you booted from it). Press Return and the next page will ask if this is an install or an upgrade. We want install. You can choose the upgrade option if you already have an old version of Red Hat...

aside a fixed amount of space for it. The downside of this is that it is much slower to access (the system has to "find" the file on the partition, it could get fragmented, etc.) and is generally foolishly dangerous. If something unfortunate happens to your virtual memory file, the system will crash and you could screw up all sorts of things.

If anything happens to Linux's swap partition, similar things may happen, but as it is on a partition of its own, this is less likely to happen and it's less disastrous if it does.

The vital question is how big to make it. This is a matter of hot debate among Linux experts. Any given 10 of them will give you at least a dozen answers, and it's probably only a matter of time before Sybex release a book entitled *Learn to Create A Properly-Sized Swap Partition in 21 Days*.

The problem is that, broadly speaking, as the swap partition grows, the less efficient it becomes. However, if it isn't big enough, efficiency isn't really a concern because you won't be able to run the things you want to. If you are chronically short of drive space, you might consider a 32Mb swap partition. Otherwise, we'd recommend one at least as large as your physical memory, or up to 2.5 times this size (exact sizes may vary depending on the geometry of your drive). If you disagree, make one whatever size you like, and see if we care...

THE ROOT PARTITION

The root partition "/" is where all the other directories and files for your Linux system are stored. This is the mother of all directories and is the only other partition required. If you don't make any of the other partitions mentioned here (apart from the swap partition) their directories will be created on this partition anyway. You should make this one as big as possible, unless you want to install any of the partitions mentioned below.

THE /BOOT PARTITION

Normally, the /boot directory lives on the root partition and is quite happy there, thank you. However, there may be a good reason for making this a partition on its own. This is the directory where the kernel image is stored and this is the bit that *LILO* loads when you try and boot up, and therein lies the problem.

LILO has a limit of how far ahead it can find a kernel image. If you have a very large hard drive, you may need to put the /boot directory in a partition of its own to make sure the kernel image stays within the range accessible by *LILO*. The exact range depends on the geometry of your drive, but if it is over 4Gb, it's well worth making a /boot partition as early as you can.

If, after installing Linux, *LILO* freezes on bootup, this is the most likely cause!

The kernel isn't that big and you can easily get away with an 8Mb partition for /boot, which will even give you some room for a few experimental kernels.

THE /HOME PARTITION

This is a useful thing to create if you are fond of experimenting. By taking this directory out as a partition, it means all of your user personal preference files and personal documents, utilities, etc, stored in their home directories are all in this partition. If you want to completely re-install Linux later (or install a different version), all these files can be saved and used by the new system, providing you don't format it during the new install process anyway...

The /home partition just needs to be as big as the files you are likely to keep there. Exact size depends on how many users you'll have, and what sort of stuff you'll be creating. If you want a partition like this, even one as small as 50Mb is useful, if only to store your email.

Linux Info

The installer options are controlled through your keyboard. Instructions appear at the bottom of the screen, but in general you should use the arrow keys, Tab and Enter to select the options on-screen.
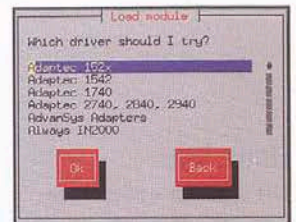
Definite Linux (and Linux in general) is pretty good at detecting hardware it can use, but you can manually specify devices too.

Step 5

Now you have the choice of what type of system to install. These choices determine which "packages" of software are installed on your computer. As the menu says, for greater flexibility, choose the Custom option. You will get the chance to specify precisely what packages to install later.

Step 6

The next thing you will be asked is if you have a SCSI adaptor. Obviously, if you do, you should answer yes. Moreover, if you are hoping to install Linux on a SCSI drive, you'll have to chose the right driver here as well. The SCSI interface is autoprobed by the installer, and most types will be set up automatically. If this fails for some reason, you
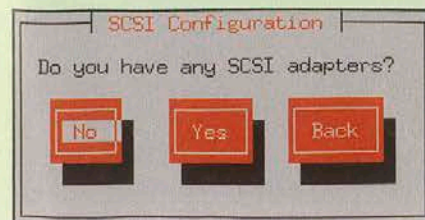

may have to specify some options here, like the I/O address of the card. You'll find more specific information on this in the Red Hat installation guide, which can be found in the doc folder on the CD.

Step 7

Now you will come to the drive setup screen. Up until now, none of your choices have actually made any physical changes to your drive. Here though we will be creating partition data. If you reset your machine at this stage, or there is a power cut, you may do some damage to your drive, so be careful!

You now get the option to use one of two utilities to set up your drive, *Disk Druid* or *fdisk*. *Disk Druid* is easier to use and understand, but won't work with drives over 8Gb.
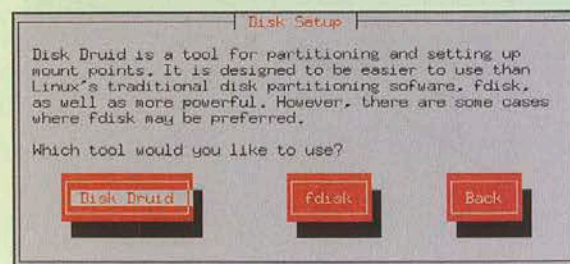
OPTION 1
USING DISK DRUID

Disk Druid is pretty easy to follow and easy to use, but if your disk has more than 1,024 cylinders, you'll have to use *Fdisk* instead. *Disk Druid* works in the same way as most of the installation menus – you can use the cursor keys and Tab to move between elements, and the Enter key to select things. There are also some shortcuts set up on the function keys, which are explained at the bottom of the screen.

Step 1

When *Disk Druid* first runs, you'll see this screen showing the current partitions on all your drives. The Device name shows the Linux-style device names for the partitions. If you only have one drive, this should be straightforward. If you have more than one drive, you'll find that partitions on the IDE primary master are labelled hdaX, where X is a number denoting the partition. The IDE Primary slave will be denoted by hdbX. SCSI drives are preceded by sd instead of hd.

Step 2
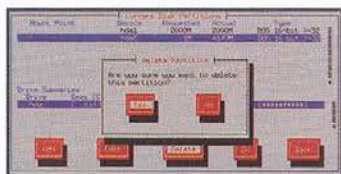
If you have a single Windows partition on a single drive, there should only be one entry on the list of drives. The Drive Summaries area of the screen will show you, among other things, the amount of unused space on your drive, and therefore how much is available to Linux.

If there is more than one partition entry and no free space, it's likely that you haven't deleted the partition you created with *FIPS*. You can do that now by selecting it and choosing Delete, but be careful that you have selected the right one! Never delete the lowest numbered Dos/Windows partition on a dual boot system, as this is normally where Windows boots from.

Step 3

The next thing we need to do is create the actual partitions. Read the main text on partition types if you haven't already done so, and the decide exactly which partitions it is that you want to create.

To create a partition, first choose Add. A screen will come up which gives you the opportunity to define the size and type of this partition.

Let's create a Swap partition. Press Tab until the "Type" section is highlighted, and then use the up and down cursors to select a Linux Swap partition.

Step 4
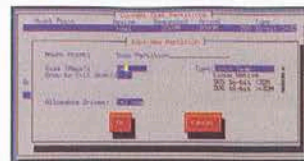
Now Tab again to the size box, and type in the size that you want for Swap space. You don't have to enter a mount point for this partition because Linux will automatically mount it as a swap drive. If you have more than one drive with space on, make sure that in the "allowable Drives" section, there is only an asterisk next to the drive you want to use (turn the asterisk on and off using the spacebar). Select OK and you will return to the main screen, but there will be a new partition listed. Hurrah!

Step 5

Repeat steps 3 and 4 for all of the other partitions you want to create. You can just create one Linux-native one with a mount point of "/" and select the "Grow to fill disk" option to finish your drive setup, but if you want more specific partitions, enter them one by one, specify the size and mountpoint, and then make sure they are "Linux native".

Step 6
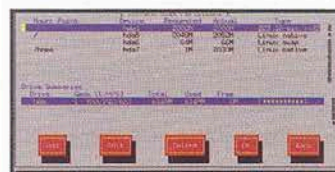
You will notice that the partitions you are creating may be shuffled around in a different order to they way you created them. This is just *Disk Druid* trying to make best use of the space, and won't affect the way the drives are mounted. If you make a mistake, you can choose Edit to change the mount point entries on listed partitions, or delete the ones you have just created without fear.

Step 7
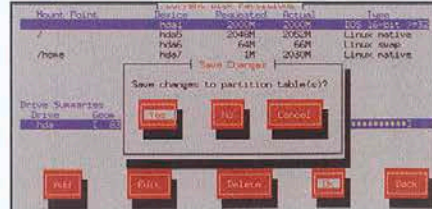
When you are happy with the way the drive is set up, choose OK or press F12. *Disk Druid* will warn you that it is about to write the data to the drive. Make sure you're happy with the way it is set up before saying yes.

Step 8
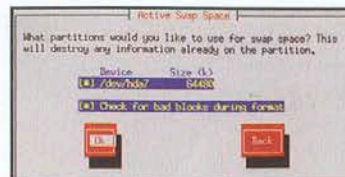
You will now be asked to format your partitions. If this is your first Linux install, you must format the drive or all sorts of problems may plague you until the end of time. Ahem. If you have a /home partition from a previous Linux install, you can leave this unformatted and be able to access all the data you stored on this path.

OPTION 2
USING FDISK

Fdisk can be a little intimidating for novice users. It is entirely based on the command line and requires you to press keys to perform actions. A complete list of the actions is shown below, though you won't have to use many of them.

If you do get confused or worried, just type q to quit from the program. No data is written to the partition until you specifically ask it to be, so don't worry about the possibility that you will damage your drive.

Use the p command to list your partitions. The numbers you'll see will depend upon the size and configuration of your hard drive. You now need to create some new partitions for Linux. Type n for 'new' and you will be asked whether this will be an extended partition or a primary partition.
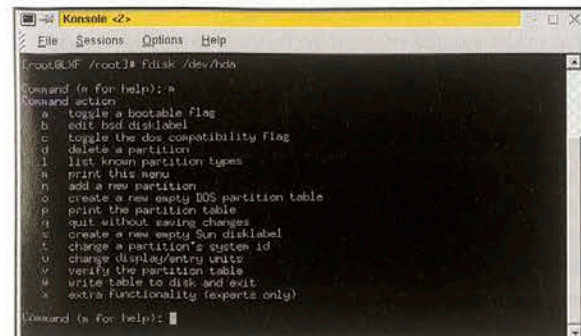
To create more partitions, you need to create a special "extended" partition, which is made to contain sub-partitions called "logical" partitions. Here we will create an extended partition with e and call it partition number 2 (partition 1 contains Windows). The first cylinder of a new partition should always be the first available cylinder, so type in a suitable value (ie, one more than the last cylinder on the existing partition). You'll then want to set the last cylinder to the last available cylinder. Again, this will depend on the size of your drive, but the number of cylinders on your drive is displayed in the line of information at the top when you enter the p command.

CREATE THE LOGICAL PARTITIONS

Now you need to create some partitions for Linux. Type n for a new partition. *Fdisk* asks whether you want to create a primary or logical partition. Type l for 'logical'.

Again, *fdisk* asks for starting and ending cylinders. It is best to create your swap partition first. Choose the first available cylinder (ie, the first cylinder of the Extended partition you created), and for a size enter +64M (or whatever size you want the swap partition to be).

Type p again and you will be able to see the new partition. You will also see that the one we just created is listed as a Linux-native partition, which is not the type we need.

FDISK COMMANDS

Command	Action
a	Toggle a bootable flag
b	Edit bsd disklabel
c	Toggle the DOS Compatibility flag
d	Delete a partition
l	List known partition types
m	Print this menu
n	Add a new partition
o	Create new empty DOS partition table
p	Print the partition table
q	Quit without saving changes
s	Create a new, empty Sun disklabel
t	Change a partition's system id
u	Change display/entry units
v	Verify the partition table
w	Write table to disk and exit
x	Extra functionality (experts only)
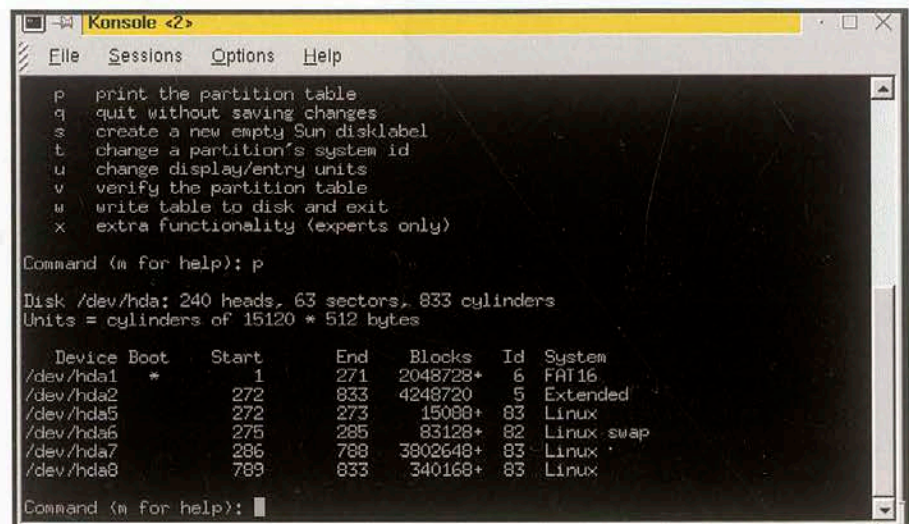
Enter the t command and choose the number of the partition you just created. *Fdisk* than asks for the hex code for the partition type. You can enter L to see a list, but 82 is the magic number for a Linux swap partition. Enter this and type p to see the changes. Create the other partitions in the same way until you are happy. When you are completely satisfied, you should use the w command in order to save the partition table.

You will then be returned to the Red Hat setup partition list, where you may choose to edit other drives or you may make changes to this one again. When you are finished, select Done, and you may be asked to reboot your computer at this time.

No data is written to the partition until you ask it, so don't worry about damage to your drive.

Partitioning in progress – nothing is written to the drive until you save.

You can get a list of the commands available by pressing 'm' and Return.
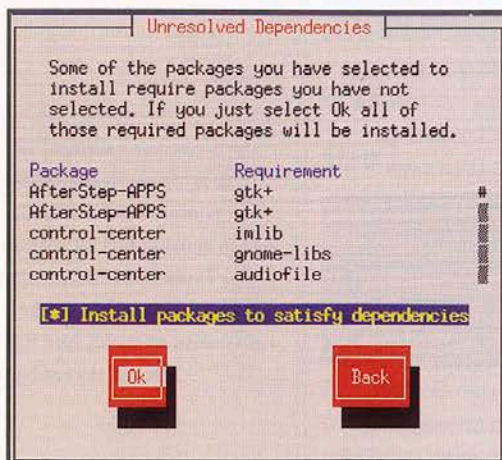
INSTALLING PART 2

infoinux Info

If your BIOS includes Power management, you may want to check the settings or even turn it off. The installation can take a long time and will not work if your computer puts itself into standby mode while the install is taking place. Alternatively, just press a key (for example, the cursor key) every now and then to keep the system awake.

Having successfully set up your partitions, the next stage is to install the software. If you chose one of the basic setups earlier, your choice here will be limited to groups of applications (such as *KDE*). If you chose the Custom setup as we advised, you'll actually be able to choose specific applications. In either case, some default applications will have been chosen for you, and if you are just trying out Linux and don't know what any of the packages actually are, you can just skip this step and move on to the section on the install process.

If you are choosing groups, select each component you wish to install and press Space. Selecting Everything (which can be found at the end of the component list) installs all packages included with Definite Linux and this will require nearly 1Gb of free disk space.

After selecting the groups, you can then select (or deselect) individual packages. Tab to the field which says "Select individual packages" and press space to highlight it, then choose OK. You'll now have a tree of the package selections. Using the arrow keys, select a group to examine and press Enter or Space. The installation program presents a list of the packages in that group, which you can select or deselect by using the arrow keys to highlight a package, and pressing Space. Some packages are highlighted in a sort of double-strike font. These are the base applications which are always installed – you can't deselect them.

In the tree, the following symbols are displayed next to the groups:

[o] – Shows that at least one of the packages in that component group has been selected.

[*] –Shows that all the packages of a component group have been selected.

[-] – Shows that one of the components is selected.

You can get a detailed description of the currently-highlighted package by pressing F1. A dialogue box will appear on the right, containing a brief description of the package. When you're done reading the description, press OK and the box will then disappear.

DEPENDENCIES

Many software packages depend on other software packages or libraries that must be installed on your system. For example, many of the graphical Red Hat system administration tools require the *python* and *pythonlib* packages. To make sure your system has all the packages it needs, the

The installer will automatically sort out dependencies.

You can choose general groups or specific packages.
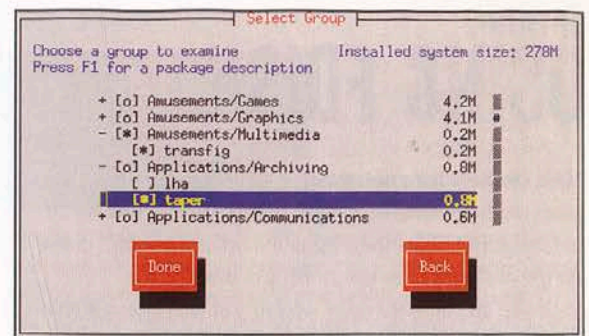
Selecting packages individually lets you install just what you want – handy if disk space is tight.

installer checks these package dependencies each time you install or remove software packages.

After you have finished selecting packages to install, the installation program checks the list of selected packages for dependencies. If any package requires another package which you have not selected to install, the program presents a list of these unresolved dependencies and gives you the opportunity to add the required packages. If you simply press OK, the program will sort out the dependencies automatically.

THE INSTALL PROCESS

After all package dependencies have been resolved, the installation program presents a dialogue box telling you that a log file containing a list of all packages installed will be written to /tmp/install.log on your Red Hat Linux system. Select OK and press Space to continue.

At this point, the installation program will format every partition you selected for formatting. This will take a few minutes depending on the size of your drive, so don't get bored and reset your computer! Just let the installer do it's work.

Once all partitions have been formatted, the installation program starts to install packages. A window called` `Install Status" will appear and give you a running commentary on the process. Several bits of information will be displayed on-screen:

Package – The name of the package currently being installed.
Size – The size of the package (in kilobytes).
Summary – A short description of the package.
Package Installation Progress Bar – A bar showing how complete the current package installation is.
Statistics Section – This section has three rows labelled ``Total,' ``Completed,', and ``Remaining". As you might guess, these rows contain statistics on the total number of packages that will be installed, statistics on the number of packages that have been completely installed and statistics on the packages that have not yet been installed.
**Overall Progress Bar** – This bar changes colour, showing how close to completion the entire installation is.

At this point, there's nothing left for you to do until all the packages have been installed. This depends on a number of factors, including the speed of your drive and how many programs you have selected. It will normally take around 20 minutes. You can amuse yourself by trying to read the descriptions of the packages as they flash by on the screen, or you can just nip off for a bit.

POST INSTALL CONFIGURATION

Now Linux is installed, but there are still lots of things to be set up and written to the various configuration files on your new system. The exact messages which appear will depend on the choices you made earlier, but we've covered most of them here. If a particular step doesn't appear, don't worry – it just doesn't apply to you.

1. Configuring your mouse
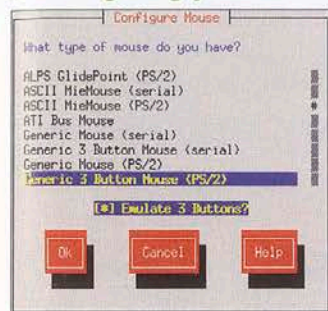
You'll want to use a mouse, so you'll have to tell Linux where to find it. Occasionally the installer may find the mouse for you and select it automatically, but usually you'll have to select it from the list like the one shown here.

If you see an exact match for your mouse, select it. If you are certain your mouse is 100% compatible with a type listed, you can select it. If your mouse isn't listed, you'll have to settle for one of the generic types – serial, 3-button serial, PS/2, etc. Hopefully you know what type of mouse yours is!

You can choose to emulate a third button here by selecting that option in the usual way. This simulates a third button whenever you press both mouse buttons together.

If you chose a serial mouse, you'll get a further screen asking you to confirm the serial port to be used.

2. Configure TCP/IP

If you are connected to some sort of network and you specified a network card earlier, you'll now have to enter your information (remember, the stuff you wrote down before?). All this information will have to be entered correctly for the network to work.
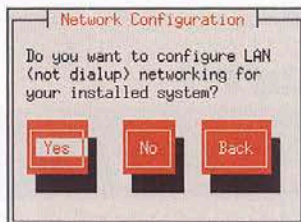
3. Configure Network

You will need to set a domain name and a host name for your system, but these can be whatever you like. If you are on a LAN, it might be sensible to use the hostname for the LAN. For example, our machines are connected to the Future network, so our domain name is futurenet.co.uk and the host name for our test machine is LXF.futurenet.co.uk. You can specify extra DNS servers here if any exist on your network.
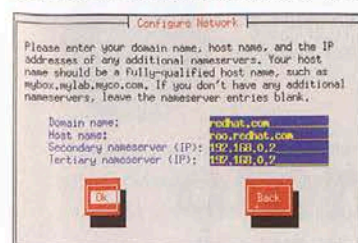
4. Time

Select your time zone from the list. Most PCs are not set to GMT so do not select this unless you know it to be true. Select the nearest City or Zone from the list.

5. Services

Next you'll see a dialogue box entitled Services. Displayed in this box is a list of services that will automatically start when you boot Linux. Scroll
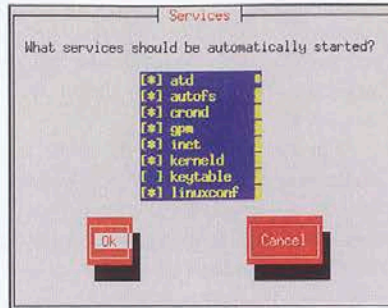
through this list and check every service that you would like automatically started every time your Linux system boots. If you're not sure what a particular service is, move the highlight to it and press F1. Feel free to play around with this, but you'd better leave keytable, atd, crond, inet, network and probably xfd selected, otherwise some things won't want to work.

6. Printers

If you have a local or network printer, you should set it up now. In fact, even if you don't have a printer, you can still set one up, and it may be useful in case you get one later. Linux will ask you for a name for the print queue and the spool directory. In most cases, it's best to accept the defaults and allow the system to search for valid printer ports.

If you have a network printer, you will have to specify the server and queue names which will obviously depend on your server. Get this information from the network administrator.

The next choice is for the printer driver. Many types of printer are supported, so you should find something close enough to yours on the list. If you have a Postscript printer, you can use the generic Postscript driver.

7. Choose Root password

You'll need to choose a password for root. The default user, root, has control over everything under Linux, so it is important that there is a password for this user. Type in something that you'll find memorable, because you don't want to get locked out of your own computer!

8. Authentication

This is an extra level of security. It doesn't matter if you don't

understand it all, just leave shadow passwords and MD5 selected. Don't select NIS unless you are on a network and there is an NIS server present. If you're not sure about an option, just don't select it.
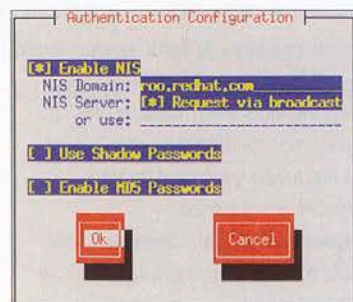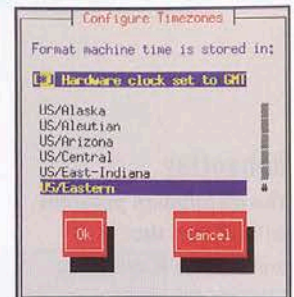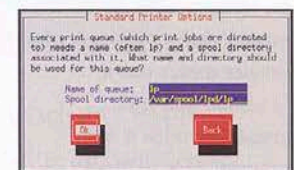

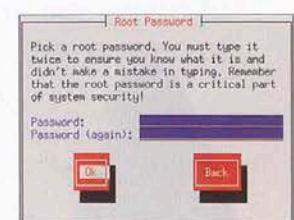*Select the time zone you want by choosing the nearest city.*
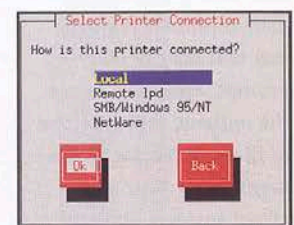
Choose the type of printer you have...
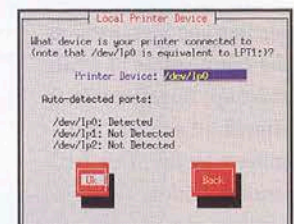
Then pick the correct port from the list...

...and create a print queue for background printing.

Make sure you pick a password you can remember easily!

Rebooting

The installation program will tell you that installation is complete. Click on OK and the system will be rebooted. At this point you should remove the CD and the boot disk if present. As your computer restarts, you will see the Lilo prompt. Press Tab to see the options, and type one of them in (Linux, if you want to try your new Linux system). You can just press Return here if Linux is the default OS.

As Linux loads, you'll see various messages as the components of Linux start up. Hopefully these should all say OK, but some may fail. This isn't usually critical. Some services may have been selected that aren't actually present (such as an NIS server), which just means that that particular service won't work.

Soon you'll be presented with a login prompt. Type in "root" as the user name, then enter the password. Hurrah! You are now a Linux user!

LILO

After installation you'll be asked if you wish to create a boot disk. For the love of everything good in the world, say YES. If you later have to re-install Windows, having a boot disk is the only easy way of booting into Linux again. Make the disk and keep it safe. Make a copy too!

Now you'll have to configure *LILO*, which is the program that loads Linux for you. *LILO* will ask you where you want to install the bootloader. Usually you'll want to install it on the Master Boot Record. Select this option and choose OK.

You will then be given the option to pass special parameters to *LILO*. You can safely ignore this unless you know what you are doing. Just click on OK to move on. Finally, *LILO* will display a list of your bootable partitions, usually one Windows one, and your new Linux one. They will have a label and one will be identified as the default. You can choose whichever you wish to be the default. This just means that, left unattended, Linux will boot into the default OS.

With *LILO* successfully installed, you will find that in future when you reboot your computer (don't do it just yet!), you will be given a choice early in the startup process of which OS to run. Pressing Return or just leaving the machine for a few seconds will select the default OS.

In order to choose a specific OS, type in the label name (you can press Tab for a list if you can't remember) and press return.
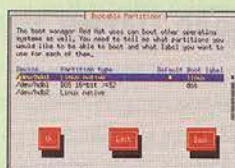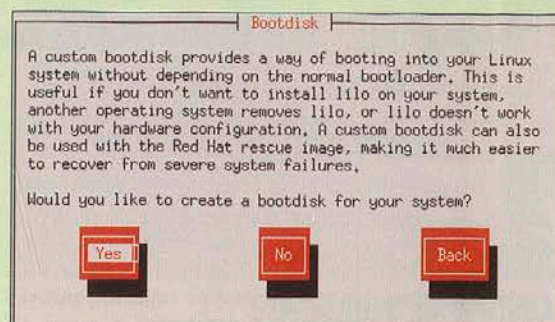
LILO *is the most common way of starting Linux on a 'dual boot' system. It's pretty easy to set up and use.*

Configuring X

X is the system used under Linux to provide a graphical interface. The X server provides a graphical interface for any programs that ask for one, so your individual applications don't have to worry about the graphics card you have. X, however, does need this information, and information on your monitor. That's right, that'll be the useful info you wrote down earlier...

Step 1

The installer will try to detect your graphics card automatically, and will highlight it on the list of supported cards. Many popular cards are listed, but if yours is not one of them, don't despair! You can choose VGA16 for unsupported cards, and this will work for any card which supports VGA. If you are sure your card is SVGA compatible, you can use this driver instead.

Step 2

The next important information you need to enter are the details of your monitor. Again, you will be given a list of supported models, but a lot of common monitors aren't listed. In this case, choose Custom and have your monitor details handy. You will need to specify the highest resolution supported by your monitor and the vertical sync range supported.

Be careful here. If you overspecify your monitor, not only will you not get a picture, but some old monitors can even be damaged if you use the wrong settings..

Step 3

Now the setup program will ask if it can probe your card. Some cards don't like being fiddled with and may completely lock up. If you select Don't Probe, you'll be asked to choose the amount of memory your card has.

Step 4

The next step is to probe for clocks. This is done after the previous probe, but some cards won't like it. If you have chosen a card from the list at the beginning, the program will warn you if the card doesn't like it. If you choose to skip the probe you will have to choose settings for the Ramdac (choose no settings if you don't know) and the modes you want X to support.

Step 5

This is the crucial test phase. Select OK and the program will start the X server. Your screen should go blank for a second, and then you'll see a grey pattern with a little X cursor, followed by a message asking if this configuration is okay. If you can't move the cursor, you selected the wrong mouse driver earlier on. This can be fixed when you have booted into Linux, but if you don't have a clue about how Linux works, it may be best to go back to the beginning of the install process.

Step 6

Finally, the setup program will ask you if you want to start X automatically. If you are setting up a desktop machine rather than a server, you may as well say yes here. This will start a graphical login every time you boot up.

What every RAD tool wants to be when it grows up
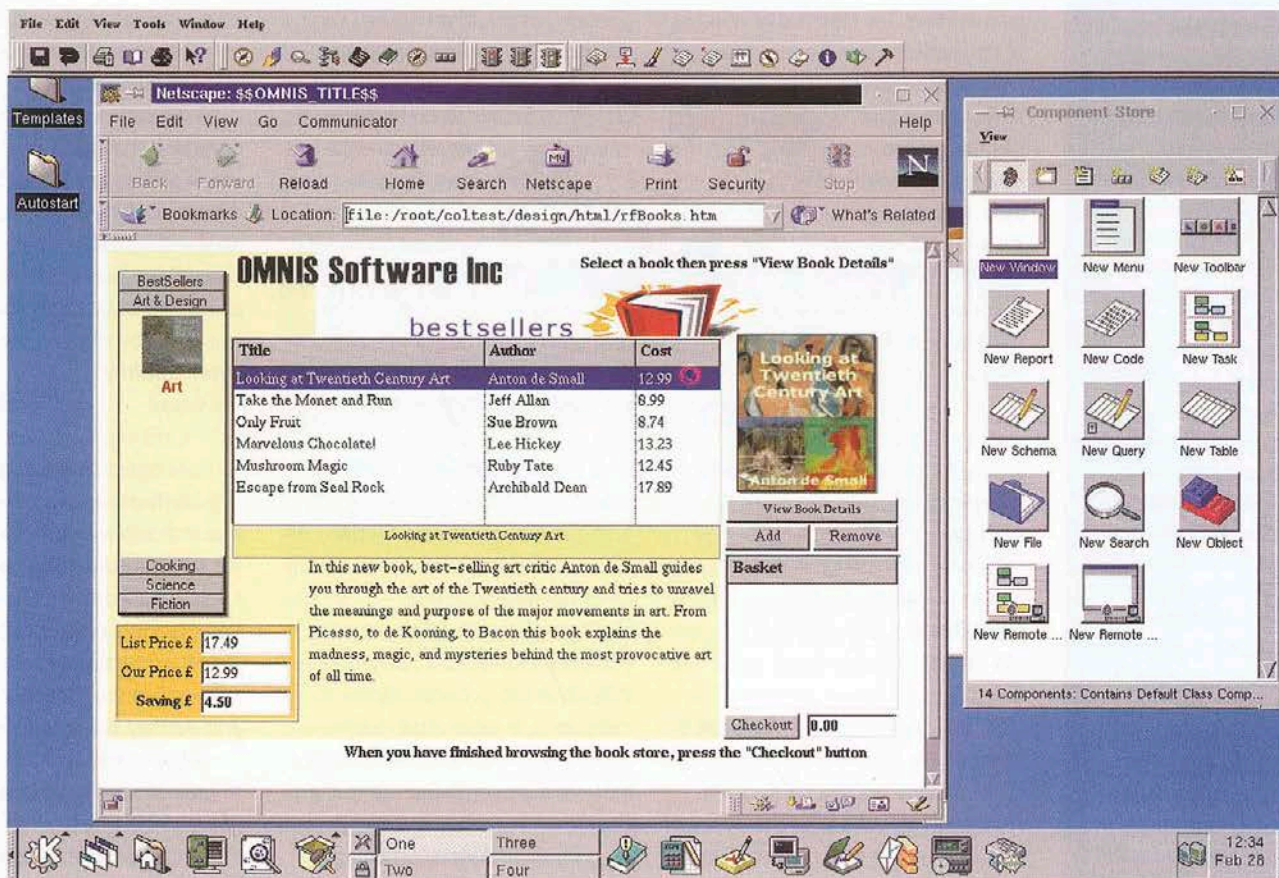
Whether you want to create an enterprise wide business solution, or an address book for your local office, you need a Rapid Application Development (RAD) tool that is flexible, powerful and easy to use. You need Omnis Studio. Whether you're on Linux, Mac, Windows or the web, you can develop where you want, deploy where you want, data store where you want.

At just £90 for the complete product, there's never been a better time to take advantage of all the functionality that Omnis Studio has to offer: drag and drop, open source components, documented APIs and powerful web functionality. Call today on 01895 453900 or download a demo from our website www.omnis.net and discover the power of Omnis Studio for yourself.

omnis.net
ahead of the curve

www.**omnis**.net

Free download of Omnis Studio demo at www.omnis.net

MAILSERVER

Share your opinions, right wrongs and demand justice by writing in to Linux Format. Drop us a line at: Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email linuxformat@futurenet.co.uk

More than Red Hat

I thought the layout of Linux Answers was excellent, but I felt that too much of the magazine was biased towards Red Hat. While I realise that Red Hat is a top-selling distribution, there are others out there as well, so maybe in each article, or at the start/end of the magazine, you could have an Red Hat, SuSE, Mandrake, etc section.

All I can say is keep up the good work that started with LA. I am looking forward to LF and have gone to my local newsagent to tell him that I want a new and exciting magazine, even though it is one that's not yet being published.
Sean Rima
via email

Well, as I hope you'll see from this issue of the mag, we are not concentrating entirely on Red Hat. However, since many of the popular distributions are based on Red Hat anyway, it is certainly the most widely used 'flavour' of Linux. Mandrake is Red Hat compatible for one. You don't actually say which distribution you are currently using yourself... Anyway, I hope you managed to seek out the mag to read my reply!

The road ahead

So Mr Veitch, Linux is the way forward, is it? The future of alternative computing and all that? My experience of Unix machines only covers the terminals at uni and I wasn't impressed, although I think they

> **Linux Answers *was a great success, but some wanted more in-depth coverage of Linux.***

probably had outdated kernels and were badly configured. I never even got round to installing LinuxAPUS on my trusty PPC powered Amiga (you really shouldn't laugh).

So, is it worth installing LinuxAPUS and the basic Red Hat distribution which appeared on an AFCD and joining the Linux revolution? I still remain to be convinced.
Colin Buchanan
c_r_buchanan@yahoo.com

Well, Linux is certainly gaining momentum as an "alternative" OS, and it is probably being installed on more computers every day than any other operating system.

If you know a bit about computers, Linux isn't really that difficult. In many cases, the installation is what most people find trickiest, so if you can master that, I'm sure you'll find the rest a doddle.

I think a PPCAmiga running Linux (especially if you have a graphics card) would be a pretty good setup. I haven't yet installed Linux on my home Amiga due to a lack of drive space, but a few years ago I ran FreeBSD on a PPC A4000 and it worked very well. There is a lot more going on in the Linux community than there has been in the Amiga world for some time now. Why not give it a try?

Penguin friendly

All computer magazines have a hardware section, normally a thinly disguised advert for the latest/fastest gubbins. However, it is much different for the Linux user as most hardware manufacturers have ignored Linux. I

would therefore suggest that your magazine should start a performance table or an article showing the ease of use of products and promoting the friendlier manufacturers.

Your readers are always considering hardware upgrades, add-ons and other changes and the most difficult problem with Linux for the newbie is that of hardware compatibility. Gradually you could build a database of UK-available equipment that would be of great use.
Peter Godfrey
via email

You will notice that we don't have any hardware reviews in this issue, and that's for a number of reasons. We will have hardware reviews in future issues. However, although a lot of hardware does work on Linux systems, most manufacturers aren't aware of this, or simply don't care and so they don't support Linux.

We don't think it would be fair to promote products which are not going to be properly supported on the Linux platform, and so, for the moment at least, we will only review hardware which is overtly Linux compatible, and which is promoted as such.

Your idea about compiling a huge database is a good one, and it's hopefully something we will be able to start creating online, but actually maintaining such a database, never mind actually creating it, really is a pretty mammoth task.

There is more than one

I know most people are running Linux to get away from Microsoft's domination of Intel contraptions, but not everybody runs Linux on an Intel box. How about some stuff for the rest

of us? I am running MkLinux on a Mac, for instance, and none of your nice software on the coverdisc will run unless you include the source! This also applies to your sister magazine (*PC Plus*) as well.
Jon Lavis
via email

A very good point. We are not going to ignore other flavours of Linux. Indeed, we'll be having a feature on running Linux on the Mac in issue 3. As far as the CD goes, we intend to supply any software in as many permutations as possible, including the source, so you should be able to use all the software included. This issue we have a full distribution of Definite Linux, which is PC-based I'm afraid, but we'll have loads of software for you next issue.

Suggestions

For the most part, I thought *Linux Answers* was excellent. It's good to have a glossy to represent our little (or not so little) niche. I have one or two comments and suggestions.

It would be nice to see a LUG (Linux User Group) announce page for LUGs which have special events, talks, etc, coming up in the near future. It would also be nice to see the Irish LUGS (ILUC, CLUG, etc) mentioned on the LUG page. Most of them, and our events, are mentioned on the **http://www.linux.ie** website.

Can we have some slightly more technical tutorials in the tutorial section too? Some of the jobs I have had trouble setting up are userspace cron jobs, log rotation and so on. Also, a brief introduction to *vi*, *emacs* and basic shell scripting would be nice. If it's not too technical, a quick intro to *LaTeX* might be good too (with *LaTeX2html* on the cover CD?).

I know it's difficult to strike a balance between getting new users using the OS and catering to the needs of the more experienced, but I think that the sooner magazines such as yourself and *PC Plus* get away from "Here's how to get Linux running"-type articles to "Here's what Linux can do" stuff – that is, getting away from

configuration and into usability, apps, etc – the better your magazine will be.

Magazines who have, for example, an install guide in every issue will inevitably lose some of the more experienced users.
Dave Neary
via email

Well, we've answered your first point, I think. The User Groups section of the mag is a regular, and any/all groups are invited to send us details of upcoming meetings for publication. If you know of a group we've missed out, send it in. We're also planning to get out and about meeting some of the groups.

As far as tutorials go, I agree

with you. There do need to be more for seasoned Linux users. This being the first issue, all our tutorial series are just starting, but further parts will include more "technical" stuff, such as shell scripting, setting up a mySQL server, programming and so on.

It is part of our remit, and indeed our duty, to welcome newcomers to the fold, so there will always be some elements of that in the magazine, but hopefully we'll find ways to do this without annoying everyone else.

Fun for all

Linux brings back the fun old days of messing about with my old Speccy and ZX81. It reminds me too of why I got into computing in the first place – the fun of hacking about with a bit of code and seeing something eventually work. Brilliant!

Are you planning on running any programming articles? It's a shame that so much space is given over to configuration and administration.

The best thing about Linux to me is the set of development tools and libraries that you get "for free". Industry standard development tools too – I've been using the likes of *make*, *yacc*, *lex*, *awk*, *perl*, *g++*, *gdb*, the X libraries, system libraries, etc since 1987.

The biggest obstacles to the average home user getting to grips with programming under Linux are probably the lack of useful documentation and either the lack of awareness that such stuff exists in the first place or the perception that it's something better left to the techies.

Either way, your magazine could provide a perfect platform for getting the message across that there is a very powerful set of development tools and

libraries in each Linux distribution just waiting to be used. Who knows what the "home hacker industry" could produce if it was aware of the potential.
Al Pearson
via email

Linux is very empowering because it does give you the freedom to tinker, and it's an excellent system to learn programming on because of the wealth of tools available.

It is an important point that, deprived of the tools to write programs (which were always supplied, even in the most basic home computers of days gone by), where are programmers going to come from? In a previous incarnation I worked on an Amiga mag where we had a "Reader Games" section. The quality of some of the submissions was excellent, and I can guarantee that all of the contributors found out a lot about computer programming.

We will be running programming features and tutorials

in the mag, hopefully something everyone will find useful.

And the winner is...

So lets get this right: DOS is dead, Windows stinks, Bill Gates, well, 'nuff said. I think I'll go and turn my machine into a Linux box.

First, using DOS, create a boot disk (one up to Bill Gates). Then I have to install my Windows lookalike interface (that's two to Bill Gates).

My machine will now never crash, as long as I regularly remake the kernel, download monthly updates, etc. Now didn't Bill Gates get slagged for service packs and new versions? Three to Bill Gates And as for command lines (DOS is alive and well) and CDs that don't automatically mount?

I don't know that a lot of the people I work with will be going down that track. I will be making some space on my machine, just so I can have a look and get the Linux experience.
Mike Roberts
via email

Hmm. I think I see what you're getting at, but:
1. You can boot up directly from the CD. If your PC won't boot from CD, how do you re-install Windows?
2. You don't need to have a "Windows-like" interface – you can use any or all of the window managers available.
3. Kernel updates, etc, are all FREE downloads, not £70 upgrades.
4. Microsoft did not invent command-line interfaces, or GUIs for that matter.
5. You can auto-mount and unmount CDs using supermount. Linux isn't for everyone, but I don't see how you can criticise it both for being too similar to Windows and too dissimilar to Window. **LXF**

WHAT WE WANT:
- Letters about the magazine or Linux in General
- Constructive criticism
- Your Opinions
- Concise points about relevant subjects

WHAT WE DON'T WANT:
- Technical question – direct those to our Q&A pages!
- Random abuse
- Nonsense rants
- 200 pages of meandering diatribe

TUTORIALS

Need some help? Our tutorials are here to help you get the most out of your Linux machine. **LINUX Format** is your guide...

Whether you are new to Linux or a weathered veteran, we hope that you are going to find some useful, informative and easy to follow tutorials here every issue.

As this is our very first issue, we're starting off with some of the basics, plus a quite useful guide to setting up your Linux box on an AppleTalk network. If you have a burning desire for a particular tutorial, write in and tell us!

Richard Drummond

IN THIS ISSUE...

USING SHELL AS A FUNCTIONAL PART OF YOUR WORKING ENVIRONMENT

Bashing the SHELL

We're not talking by the seashore here, were talking how to use the most powerful part of your Linux system, the shell. Window managers are all well and good, but sometimes you just need the raw power and flexibility of the shell – if you're not using it, you're not really running Linux.

p63 Our shell supremo Dave Coulson is your guide to the first part of this series, which begins with some simple and some not so simple commands.
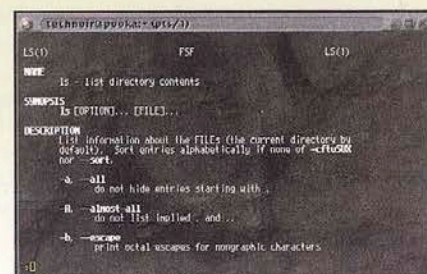
A shell window - your gateway to power, performance and flexibility

HOW TO FOOL A MAC NETWORK INTO ACCEPTING A LINUX BOX

APPLE networking

Not everyone in the world can sure the pure beauty of a Linux network, and Linux boxes hate to be alone, so sometimes the only solution is to adopt subterfuge.

All Apple Macs have networking capability built in, and in fact, it's actually a very good networking system, so if you

p67 live in a world full of these computers (you might even be running one), you might care to join us as we probe the depths of AppleTalk with Anders Brownworth.
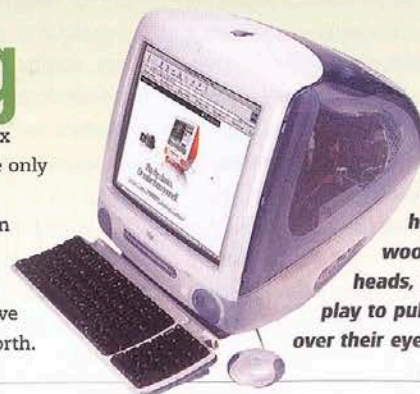
He-he-he-he- so much wool in their heads, it's child's play to pull it down over their eyes.

GETTING YOUR OWN WEB SERVER UP AND RUNNING

APACHE rides again

Have you ever wondered why Apache is the most popular web server on the planet (by a very large margin). Well, of course, partly because it is free, partly because it runs on Linux, but probably more importantly because it is robust and pretty easy to set up.

If you don't believe us, why not join Joel Rowbotham as he

p71 explains how you too could be powered by Apache! We'll have more Apache-related tutorials to follow in the months to come, so keep reading!

With Apache you can easily create your own website for all occassions.

Bashing the SHELL

The Linux shell is more than just a command line. **David Coulson** shows you how to use the shell as a functional part of your working environment.

As soon as anyone mentions using a shell, command line, or script, most people tend to run a mile. They think of it as a nasty way to do something when the GUI fails, or as a last resort when their machine dies. However, with Linux, the shell becomes an integral part of your working environment. You launch applications from it, access remote machines, administrate your box and you can even configure your X Window setup via it.

There are a wide variety of shells available for Linux, all with their unique syntax and commands, but they all use the same commands for general usage, simply because 'ls' is actually a binary, rather than a shell command. We're only going to focus on Bash here as it's the default shell on almost all Linux distributions, although almost all of the basics apply to each of the various shells, and are used in the same way.

THE BASICS

The first thing people will want to do with their shell is to navigate their file system. Remember, with Linux there's none of that A:\, C:\, D:\ faff, it's all mounted somewhere on /, so it's far easier to organise your devices and mounts. You could quite easily be fooled by trying to use DOS style commands, as common ones such as 'cd' and 'dir' do exactly the same on both platforms.

The main file system commands are fairly self explanatory, until you start doing complicated things with them. 'cd', rather unsurprisingly, lets you change the current working directory of the shell. In DOS you might do 'cd ..', to move back a directory. You do exactly the same in Linux, although it's worth remembering that '..' and '.' are actually just directories, which refer to the parent director and current directory respectively. Here, we hit the snag of filenames in Linux. Unix type platforms have been able to use long files for years, so there have been allowances made for filenames with strange characters in.

For example, if we had a directory called 'Test Code' and we tried to cd into it using 'cd Test Code', it would fail, as cd would

think you're trying to get into the 'Test' directory. The simplest way to get around this is by putting the directory name within quotes, 'cd "Test Code" ', which causes the shell to send the cd command 'Test Code' as a whole argument, rather than 2 separate arguments. Another option is to escape the space character. Escaping characters is quite an important concept, especially if you accidentally put a quote in a filename.

Using this method, you'd issue the command

```
cd Test\ Code
```

to the shell. As far as the 'cd' command is concerned, it's getting the string 'Test Code' as it's first argument in both cases, since all the processing of your shell command is done by bash. The 'ls' command is almost comparable to 'dir' in DOS, so most people won't find it →

> **Unix-type platforms have been able to use long filenames for years...**

Your shell may look a little different to this one, but the commands should work just the same.

Useful Commands.

cd - *Change current working directory*

pwd - *Displays current working directory*

ls - *Lists contents of current directory*

mv - *Moves, or renames, a file or group of files*

cp - *Copies files from one path to another*

rm - *Deletes files*

whoami - *Tells you who your logged in as*

ps - *Lists current processes on your shell*

kill - *Sends a signal to a process, usually in order to kill it*

grep - *Only outputs the parts of the input which match a particular pattern*

more - *Displays input content a page at a time*

man - *Displays manual page for a command*

info - *Displays the info page for a command*

w - *Lists all users who are logged in and their terminal*

wc - *Displays the number of characters, words and lines in the input or a file*

➔ awfully confusing. 'ls' on it's own, simply lists the files in the current working directory, which on it's own isn't exactly exciting. However, it will do simple searching and sorting. You might have a directory full of images and you want to see what the newest jpeg images are called and when you created them. The first thing to do is to only pick out the images ending in '.jpg'. Unlike DOS, Linux shells don't care about the period in the filename, so '*.*' doesn't refer to every file, just those which contain a period. The command 'ls *.jpg' would list all files ending in '.jpg' in the current directory.

As with escaped characters or quoted strings, wildcards are transferred into the matching filenames by the shell, so the ls command actually gets a huge list of files in the current directory which match your pattern. Just like a space, if you have an asterisk in a filename, pop a \ before it, or pop the filename in quotes, and the shell won't try to match the pattern. If you want to sort the filenames by time, you use the '-t' switch on ls. You can put the '-t' before or after your pattern or filename; ls doesn't care in the least.

To have it display the time of last modification, the '-l' switch is used, resulting in a

```
ls -t -l *.jpg
```

command. With the standard GNU tools, which almost all general Linux commands are, you can combine switches beginning with a single hyphen into a group, so

```
ls -tl *.jpg
```

is equally valid. Switches which begin with 2 hyphens, such as '--reverse' in the case of ls, are handles individually as they have a longer, more user friendly name, usually accompanied by a value, as in the case of '--sort=t'. But what if you wanted to list both jpegs and pngs using ls? As the shell simply creates a list of filenames, which is then passed to ls, there's nothing stopping you simply making the shell create a long list of filenames.

```
ls -tl *.jpg *.png
```

would work exactly as you'd expect, listing both jpegs and pngs by time. With switches, you hit a small snag; What if your file begins with a hyphen? If you escape it or pop it within quotes, it gets passed to the command as if they weren't there at all. Trying to delete a file '-foo' with the command

```
rm -fr "-foo"
```

grep is a very useful command to get to know, you may not use it much on it's own, it's real power is when it is combined with other commands.

would cause rm to bail out because it doesn't understand the '-o' switch. The proper way for commands to handle this situation is to assume that everything which comes after ' -- ' is not a switch, so 'rm -fr -- -foo' would delete the '-foo' file. This same style of filename wildcards applies to any command you type into the shell.

If you wanted to display the contents of all your .c or .h files, you could use

```
cat *.c *.h
```

The shell would expand on the wildcard in the order it finds them in the directory, so you'd end up seeing all the .c files in alphabetical order, then all the .h files sorted the same way. However, as a resultant of this, commands such as 'mv' might not work how you want them to.

By trying to rename all your .txt files to .text using

```
mv *.txt *.text
```

which at a first glance might look correct, you are in fact trying to copy all the *.txt and all but the last *.text into the directory given by the filename of the last .text file. It may sound confusing, but if you think about what 'mv' is actually being told to do, it makes perfect sense. Another nifty way of expanding filenames, although not actually as part of a command line, is to use the 'tab' key.

By simply typing in the first few characters of a filename, pressing tab will attempt to auto complete the filename. If you've got more than one file with similar filenames, pressing tab twice will give you a list of these files, so you can just add an extra character to differentiate between them, then press tab again to auto-complete. The auto-completion does not differentiate between files and directories - all valid pathnames will be potentially completed using Tab.

THE NOT-QUITE-SO BASICS; PIPES, PROCESSES AND REDIRECTION..

Earlier I mentioned that 'cat' displays the contents of files. Not entirely accurate. From a shell point of view, it writes the contents of the files passed to it into stdout. It just so happens that as default, stdout in a shell is the terminal display (be it Linux console, xterm, etc). 'ls' outputs it's data in exactly the same way. So, this gives us the opportunity to send the data somewhere else. The easiest way to do this is via redirection. To send the output of 'ls' to a particular file, say ls.txt, we would use

```
ls > ls.txt
```

Only the part of the line up to the '>' is actually sent to 'ls', so it doesn't know anything has changed. You can also send data into a program using the '<' character. For example, if you wanted to read a text file a page at a time you could use 'more < file.txt'. But what if we wanted the do something to the output from one command using another? We could use a temporary file and send the output from the first there, then run the second command using that file as it's input, but that's both messy and tedious. Pipes let you send the output from one command into another, as it's input. So if you had a directory with a lot of files in you could use 'ls -tl | more' to view them a page at a time.

There's virtually no limit to the number of pipes you can have in a command line, so you can quite easily string together a line containing 6 or 7 commands, so you get only the information you're after. Yet another option to use, especially if you've got 2 complicated programs running on each side of the pipe is to use a named pipe. This is just a special file on the

man pages will give a useful, if terse overview of a command. The info pages are much more descriptive.

disk which you can connect programs to, then pipe data into it. To create your named pipe use the 'mkfifo' command. You could create a pipe called 'ls-pipe' using the command

```
mkfifo ~/ls-pipe
```

then connect grep to it using

```
grep "^-" < ~/ls-pipe
```

which will only output files, not directories or links. To put data into your named pipe, it's just a matter of redirecting output from 'ls -tl' to the pipe as if it were a file.

One thing to remember when using this type of pipe is that the command which is sending data to the pipe will not close until the receiving program has collected all the data. The receiving program will also automatically close when the input has finished. Unfortunately, you can't connect multiple programs to the output, as you might expect, so processing the input in 2 different ways requires you to create 2 pipes then do 'ls -l > ~/ls-pipe > ~/other-pipe', but you can use the '>' and the '|' redirection operators within the same command, as the stdout of the program is sent to both.

BACKGROUND TASKS

But, what happens when I type 'ls'? As mentioned earlier, it's a binary program, rather than being internally handled by the shell as 'cd' is. When you run ls, the shells spawns a process called 'ls', which then performs the task before finally shutting down, and in the process stops you from running any other programs from the command line.

With something simple like 'ls', not being able to do anything while it's running isn't much of a problem, since it's practically instantaneous. However, once you run more involved programs from the shell, *xv* for example, having a 'dead' terminal while you use the program is a bit annoying. If you've got *xv* up and running, you don't have to close it down to be able to use your terminal again. Simply press Ctrl-Z and you'll get your shell back. Not a lot of use, you may think, as it's stopped *xv* working. Ctrl-Z is actually a process management command which suspends the current task, so it's up to you to get it back up and running again. You can either have it run in the foreground again or, as we want, in →

> ## You can't connect two programs to the output of a pipe, as you might wish – you'll need another pipe

```
technoir@pooka:~ (pts/1)

LS(1)                    FSF                    LS(1)

NAME
     ls - list directory contents

SYNOPSIS
     ls [OPTION]... [FILE]...

DESCRIPTION
     List information about the FILEs (the current directory by
     default).  Sort entries alphabetically if none of -cftuSUX
     nor --sort.

     -a, --all
          do not hide entries starting with .

     -A, --almost-all
          do not list implied . and ..

     -b, --escape
          print octal escapes for nongraphic characters
:0
```

→ the background. Simply type 'bg' and press return, and *xv* comes back to life and your terminal is still working. It works, but it's a bit tedious to do that every time you want a program to run in the background, so the shell provides a work around. Just add a '&' to the end of your command and hey-presto, it's in the background automatically. The sharp eyed among you will have noticed that it spat out a couple of numbers in the form '[x] xxxx'. The first number is simply the job number for your


```
technoir@pooka:~ (pts/1)
[technoir@pooka technoir]$ ps
  PID TTY          TIME CMD
12259 pts/1    00:00:00 bash
12856 pts/1    00:00:00 ps
[technoir@pooka technoir]$ ps -a | grep "x$"
 9743 tty4     00:00:00 startx
[technoir@pooka technoir]$ []
```

Grep to the rescue once again, as in this case we parse all the running processes for one which ends in 'x'.

command in the current shell, so if this is the first background command running in that shell, it's say '[1]'. The second number is far more important, as it's the process number for your program on the system. If you wanted to kill off a process currently running in the foreground, pressing 'Ctrl-C' would kill it off, but with an application running in the background, you don't have that ability. So, you can either bring it to the foreground, using 'fg', or kill it via it's process number.

The 'kill' command is used to send particular signals to processes, such as SIGTERM, which kills most processes. 'kill xxxx', where 'xxxx' is the process number, will do the job for you. Again, we hit another issue; What if you've forgotten the process id for your application? The 'ps' command provides you with a list of processes which are running on the current

terminal, so by typing 'ps' at a new command prompt, you'd be told it's running 'bash' and 'ps', along with their process ids and their current CPU time usage. Even from a different terminal you can see your current processes by adding the '-a' switch to your ps command. Unfortunately, if your running lots of different programs you'll have difficulty hunting through the list for the one your after. As usual, pipe comes to the rescue. The GNU tools provide you with a rather handy suite of applications for string manipulation and the one we're interested in right now is 'grep', which is an acronym for 'Global Regular Expression and

As usual, pipe comes to the rescue, teaming up with the grep command...

Print'. We need not concern ourselves with the actual meaning of that, but it basically means "Read stdin and only output lines which match our pattern". In the case of our ps output, we only want lines which end with the name of our process.

```
ps -a | grep "xv$"
```

is the command we're after. The '$' symbol at the end of the regular expression means 'end of line', so rather than matching lines which contain *xv* at any point, we restrict it to those which end with *xv*.

Reusing old commands is always a nice idea, so the bash shell provides you with a history of past commands. The simplest way to use this is by using the up/down cursor keys to cycle through previous commands and at any point you can either press return to execute that command, or use the left/right arrows to edit it. If your after a command which you used in a previous session, you can use the '!' to access them. You simply type '!', then the first section of the command your looking for, and it will execute the last command you used which began with those particular characters.

WHAT TO DO IF YOU GET A BIT STUCK

Trying to figure out exactly what you can do with a particular command is usually pretty tedious, but even something as simple as 'ls' or 'pwd' have their own manual pages. Rather unsurprisingly, the command 'man' is used to gain access to these by passing the command your interested in as an argument to man. Another way is to use the built in help provided by the command. Usually this is nothing more than a list of valid switches, but it's better than nothing. GNU apps usually use '--help' as their help switch, but others may use '-help' or just '-h'. Yet another way is via 'info' pages. These are more organised than the man pages and give more involved descriptions and examples showing how to use the command. Once you've understood this lot, you'll be ready for our scripting tutorial in next issue! **LXF**

TIPS AND TRICKS

1) Unsure about a particular command line option? Most programs have help available when you pass '--help' as a switch
2) You can include more than one command on the same line, separating them by a ';'.
3) User's home directories are refereed by ~user-name/, and the current users home directory is under ~/
4) Looking for that misplaced file? locate

uses the daily filesystem database build to search for files containing the pattern you pass to it.'locate foo' would list all files with foo in their path.
5) ls isn't particularly helpful with default settings, so add '--color=tty' to it and you get pretty colours, depicting the type of file. Simply add an alias to /etc/profile, alias ls='ls --color=auto', and it'll do it all the time.

Connecting to an
APPLE MAC
NETWORK

Anders Brownworth explains how you can fool a Mac network into accepting a Linux box on friendly terms...

You've spent all this energy getting a Linux machine running, but you still have a network full of Macs. Wouldn't it be nice to get them to play together nicely? Well, you're in luck. *Netatalk* is a program that makes your Linux machine look like a Macintosh fileserver on a network. In this article, we're going to show you how to set it all up.

Netatalk supports two types of AppleTalk: Classic AppleTalk and AppleShareIP (AppleTalk over TCP/IP). The original *Netatalk* by the University of Michigan's Research Systems Unix Group supports Classic AppleTalk. Since then, Adrian Sun has been developing *Netatalk* to include AppleShareIP and numerous other bug fixes and enhancements. We are going to install Adrian Sun's version of *Netatalk, netatalk-1.4b2+asun2.1.3.tar.gz.*

You can download the source code, a precompiled version or the Red Hat rpm

DOWNLOAD NETATALK

There are several different ways you can download *Netatalk*. You can download the source code, a precompiled version or the precompiled version as a Red Hat rpm. We'd recommend you download one of the precompiled versions. If you wish to compile *Netatalk* from the source, take a look at the Netatalk-HOWTO located at the following website:
http://thehamptons.com/anders/netatalk/.
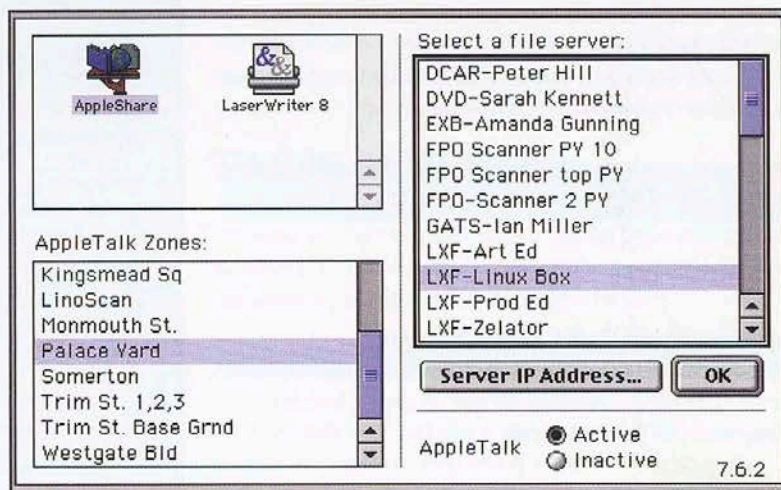
Ha! Our evil plan has worked, and our Linux box has infiltrated the network.

Original source:
ftp://ftp.cobaltnet.com/pub/users/asun/release/netatalk-1.4b2+asun2.1.3.tar.gz
Precompiled version:
http://thehamptons.com/anders/netatalk/mirror/netatalk-1.4b2+asun2.1.1.x86-shadow-linux.tar.gz
Debian .deb version: http://cgi.debian.org/www-master/debian.org/Packages/stable/net/netatalk.html
RedHat .rpm version:
ftp://contrib.redhat.com/pub/contrib/libc6/SRPMS/netatalk-1.4b2+asun2.1.3-6.src.rpm

If you're running Red Hat and grabbed the rpm, install it like this:

```
rpm -iv netatalk-1.4b2+asun2.1.3-
6.src.rpm
```

If you grabbed one of the precompiled versions, decompress and install them like this:

```
> tar -xzvf netatalk-
1.4b2+asun2.1.3.tar.gz
> cd netatalk-
1.4b2+asun2.1.3
> make install
```

CLASSIC NETWORKING

For Classic AppleTalk networking, *Netatalk* requires AppleTalk support in the kernel. Luckily, most major distributions include AppleTalk support, such as RedHat. If you don't have it, you will have to recompile your kernel. To check whether you have it, try this: dmesg | grep -i apple.

The installed version of *Netatalk* lives in: /usr/local/atalk/. Server binaries and config files are in /usr/local/atalk/etc/ and userland binaries are in /usr/local/atalk/bin/.

Red Hat and Debian may store the package in non-standard locations. You may have to do something like this:

```
find /* -name atalkd -
print
```

to discover the binary.

AppleShareIP and Classic AppleTalk Shares:

afpd, the program that makes AppleTalk shares available in both Classic AppleTalk and AppleShareIP, gets its options from the configuration file called afpd.conf.

The simplest afpd.conf file just includes a − which stands for "this machine" with all the default options. A slightly more complicated example would be a file containing one word (for example, "balboa"), as the name of the server.

Let's make this clear with a few examples. Let's assume we have a server called "balboa" with the following afpd.conf file:

```
-
Balboa2 -port 12000
"Bad Boy" -port 12001
"Out There" -port 12000 -address
206.114.89.46
```

The first line (−) gives you a server called balboa. The second line creates a server called Balboa2 on port 12000 and the third creates yet another server called Bad Boy on port 12001. The last line creates a server called Out There, bound to port 12000 on virtual IP address 206.114.89.46.

```
"Guest Volume" -nocleartxt -
loginmesg "Pay homage to my guest
                            server!"
```

This line will create a dedicated guest server. -nocleartxt disables password login, hence the guest nature of the server. A login message will be shown after the guest user connects.

```
"The Vault" -noguest -port 12000
```

With this setting we will get a volume on port 12000 called The Vault which will not allow guest access.

```
"Inner Sanctum" -notcp -defaultvol
<path> -systemvol <path>
```

Here we define a server called Inner Sanctum which does not use TCP/IP (DDP only) and uses an alternate AppleVolumes.default and AppleVolumes.system file. Don't forget to replace <path> with the correct path and file name.

APPLEVOLUMES FILES →

To establish a Classic AppleTalk share, we have to edit /usr/local/atalk/etc/AppleVolumes.system and add the share:

Setting up the configuration files will allow you to share any Linux directory on your system.
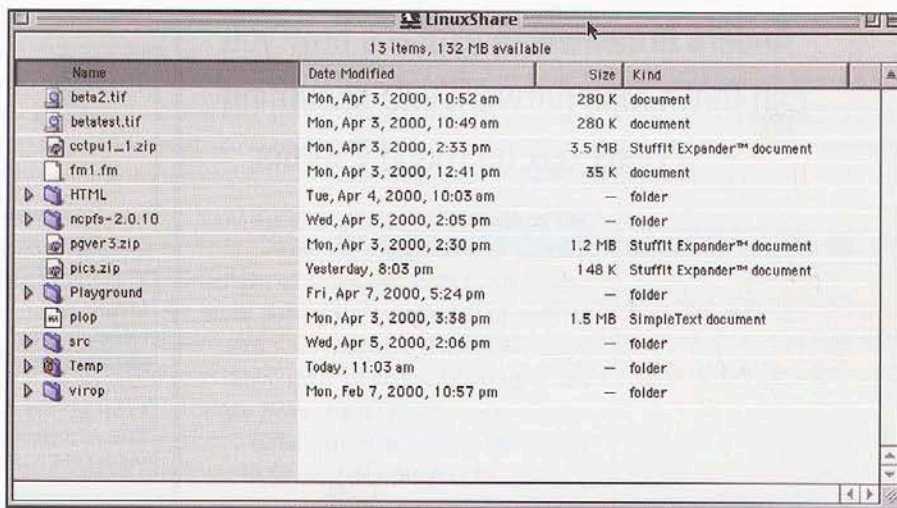
And here is a Linux directory being accessed from a Mac.

→
/data/www "Webserver"

This creates a volume called Webserver which exposes the contents of /data/www.

The other thing that this file controls is type-creator mappings. Unlike Linux files, all Macintosh files consist of two pieces, a resource fork and a data fork. The resource fork stores things like icons, file types and a mapping to the application that was used to create the file.

Because Linux files are essentially one long list of characters instead of a resource fork and a data fork, *Netatalk* writes the resource forks in files inside a directory called .AppleDouble. Therefore, a pair of Linux files represents one Macintosh file (somefile.txt and .AppleDouble/somefile.txt). Consider the following type-creator mapping that would be in afpd.conf:

.txt   TEXT   MSWD

This tells afpd to tell Macintoshes that the file type is of type TEXT and the creator is MSWD (Microsoft *Word*) when it runs into non Macintosh files ending in .txt. The server will do this for all files that do not have a .AppleDouble counterpart.

THE ATALKD DAEMON

atalkd is the Classic AppleTalk kernel interface. It will serve as the link between the kernel AppleTalk module and the rest of the classic AppleTalk functions of *Netatalk*. It will even take care of AppleTalk routing between multiple network cards.

This program is controlled by the atalkd.conf file. The simplest version of this file is an empty file which causes atalkd to overwrite it with some reasonable options. You should specify all of your network cards in this file. If you specify more than one, atalkd will route AppleTalk packets between the cards.

Consider the following examples:

eth0 -net 153-174 -addr 154.212 - zone "The Far Side"

Each line of this file specifies a network interface (eth0). The rest of the parameters are optional. In the example, the eth0 interface is on a network of AppleTalk addresses between 153 and 174. The card will be configured to use the AppleTalk address 154.212 and exist in the zone "The Far Side".

eth0 -net 153-174 -addr 154.212 - zone "The Far Side"
eth1 -net 175-200 -addr 182.318 - zone "Bloom County" -seed

The second line causes eth1 to exist on a network between 175 and 200 with the address 182.318 and seed the zone "Bloom County". Therefore, atalkd will act as a router on this network and acts as a bridge between the two networks. Note that the network numbers may not overlap, nor can there be another router on the eth1 network. If another router is discovered that disagrees with what atalkd is trying to set, atalkd will exit.

STARTING UP THE SERVER

Now comes the fun part. Try to test a startup of the *Netatalk* server by running the rc.atalk script in the /etc/init.d directory. Depending on your install, you may have to do a find to figure out where the startup script is installed.

find /* -name "*atalk*" -print

Type "./rc.atalk start" and give it a few seconds. This will start up a Classic Appletalk fileserver and Appletalk IP if you have any TCP options set in the afpd.conf file. Atalkd will take the longest to start as it checks out the network before registering itself.

Alternatively you could just run an Appletalk over TCP/IP server by having a TCP server configured in the afpd.conf file and launching afpd by hand:

/usr/local/atalk/etc/afpd -F /usr/local/atalk/etc/afpd.conf

If everything seems to have started up without complaining, go to a Mac and open the Chooser (under the Apple menu on the left-hand side). Click on AppleShare and see if your *Netatalk* server shows up. For AppleShareIP you have to click the AppleShareIP button and type your machine IP or name in.

With some luck you will be presented with a password prompt. You can only log into the server with a non root account that has a valid shell and a password of eight or less characters. Guest access is also permitted unless denied in afpd.conf. If you are able to login, you will get a list of volumes which you can

You will have to edit the config files manually to set the network up properly.

```
(Konsole)    (Konsole <2>)    (Konsole <3>)    (Netscape: Welcom...    /etc

/etc/atalk/AppleVolumes.default

File  Edit  Options  Help

# This file looks empty when viewed with "vi". In fact, there is one
# '*', so users with no AppleVolumes file in their home directory get
# their home directory by default.
#
# volume format:
# path [name] [casefold=x] [codepage=y] [options=z,l,j] \
#    [access=a,@b,c,d] [dbpath=path] [password=p]
#
#
# casefold options:
# tolower    -> lowercases names in both directions
# toupper    -> uppercases names in both directions
# xlatelower -> client sees lowercase, server sees uppercase
# xlateupper -> client sees uppercase, server sees lowercase
#
# access format:
# user1,@group,user2  -> restricts volume to listed users/groups
#
# miscellaneous options
# prodos              -> make compatible with appleII clients.
# crlf                -> enable crlf translation for TEXT files.
# noadouble           -> don't create .AppleDouble unless a resource
#                        fork needs to be created.
#
# codepage=filename   -> load filename from nls directory.
# dbpath=path         -> store the database stuff in the following path.
# password=password   -> set a volume password (8 characters max)
```

select and mount. Use one of the startup scripts included with *Netatalk* to launch your server when your system boots. There is a variety of ways that a machine will boot up. Some have an rc.local file where you can specify the path to an rc.atalk script while others use a symbolic link with a naming scheme to determine the order in which startup scripts will run. Take a look at /etc/rc.d/ or /etc/init.d/ for startup scripts.

Note:

If you want to serve more than five connections, you must supply a max connections variable to afpd. The -c flag will set this.

```
afpd -c 25
```

OTHER THINGS YOU CAN DO

Two-Way Encrypted password authentication is supported by *Netatalk*, but to do this, the server needs to know what the clear-text password is.

Because the system passwords are encrypted, the server doesn't know the clear-text version, so you have to supply it. By creating a .passwd file for each user in their home directory with their clear-text password contained in it, Two-Way Encrypted password authentication can be established.

However, it is obviously quite insecure to leave the clear-text passwords in a predictable location, so *Netatalk* requires that each user's .passwd file be owned by that user and his default group, and read-write access to that user and to nobody else.

Simply put, you should do this to the file:

```
chown user.group
```

```
/home/user/.passwd
chmod 600 /home/user/.passwd
```

where "user" is the username and "group" is that particular user's default group.

PRINTING

Netatalk doesn't only act as a fileserver – it can also do print spooling or even print to existing AppleTalk spools. AppleTalk print spooling is accomplished with a program called *papd*. The concept is that you have a local (parallel or serial) printer that has Mac drivers and is already set up under Linux. *papd* will advertise itself on the network as a print spool and accept jobs from Macintoshes, (print jobs are PostScript files) and then *papd*
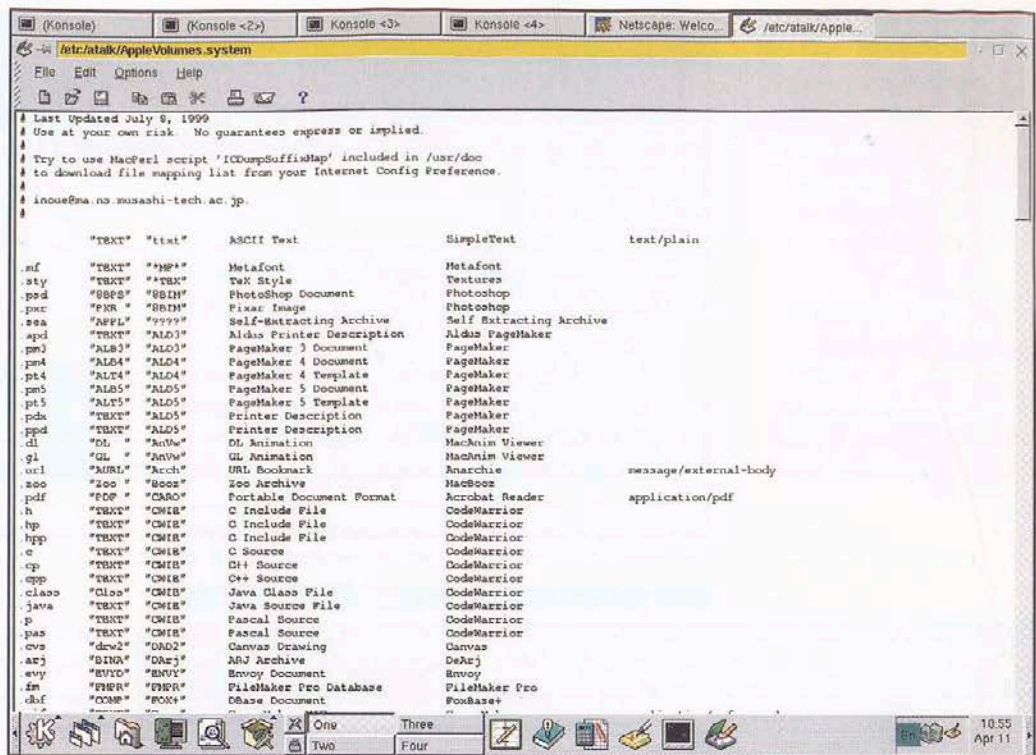
feeds this information to the Linux printing system (lpd). Examples of setting *papd* up are available on the web at:
http://www.giub.unibe.ch/~eugster/appleprint.html

The reverse of this (printing to a Macintosh printer from Linux) is done through *pap*. To see how to use *pap*, check out:

```
man pap
```

Just like there are utilities to ping computers over TCP/IP, there is "aecho" for Appletalk. It works just the same as you would guess. "getzones" will return a list of the Macintosh zones on the network, and "nbplkup" with no options will return a comprehensive list of all the AppleTalk services available on the local network.

There are many other handy little programs to discover with *Netatalk*, so don't be afraid to get your hands dirty.

ENJOY!

With some luck, you have followed us through the installation and configuration of *Netatalk*. A Linux machine running *Netatalk* and *Samba* (the Windows equivalent of *Netatalk*) gives a Windows NT machine a run for the money. Tack a webserver like *Apache* on there and you can have a fully functional industrial strength web development platform built with free software. Enjoy! **LXF**

There are many handy progams to discover with Netatalk, so don't be afraid to get your hands dirty

Macs are difficult beasts if they don't know the origins of the files that they are dealing with. This file sets up some of the common filetype associations.

APACHE RIDES AGAIN

Need more control over your own web pages? Try installing your own Apache web server, says **Joel Rowbottom**.

Unless you've had your head in the clouds for the past eight years, you'll have at least heard of – if not used – the World-Wide Web. Pages of information are displayed to users, often providing live, dynamic data, as well as static information. In order to serve this data to users, be that on the Internet itself or through a private network (an Intranet), you need to have access to a web server. This is a piece of software which serves requested pages to clients.

Almost every single Internet Service Provider (ISP) provides publicly-accessible web servers for customers' use (that's the very nature of their business, of course). Let's say, however, that you want to publish company data on a machine which is only available to users inside your own internal network, or perhaps you've got a permanent connection and would like more control over your site. That's where your own web server comes in. Under Unix-based systems it's called an "http daemon", or "httpd" for short.

The most popular web server software which will run on Linux is called *Apache*. This piece of software grew from "NSCA httpd", developed at the National Centre for Supercomputing Applications in Illinois. A number of authors who had submitted patches for the NCSA httpd server banded together to form the Apache Group, so named because it was literally "A Patchy Server".

OBTAINING APACHE

Most versions of Linux are shipped with a pre-compiled copy of *Apache*, which you may need to discard if you opted to install it when you first installed your distribution. The simple de-installation commands vary from distribution to distribution, but the most popular are:

Red Hat:
```
rpm -e "rpm -qa | grep apache"
```

Debian (and derivatives):
```
apt-common remove apache
```

Slackware (and most other distributions):
```
rm -rf /var/lib/httpd
/usr/local/apache
```

If you're planning on running a dedicated web server, it is invariably more sensible to start with a clean distribution and

The most popular web server software on Linux is Apache.

leave *Apache* out of the initial installation altogether. Also, bear in mind that some online documentation readers depend on *Apache* being installed (especially distributions based on Debian).

It's a good idea to obtain the latest stable version (at time of writing this was 1.3.12). You can grab a copy of the source distribution from **http://www.apache.org/**, although I tend to use one of the UK mirrors since it's nicer →

You're sure to come across the Apache logo on your web travels – there is a good reason for this!

The Apache Software Foundation website is a good source of information.
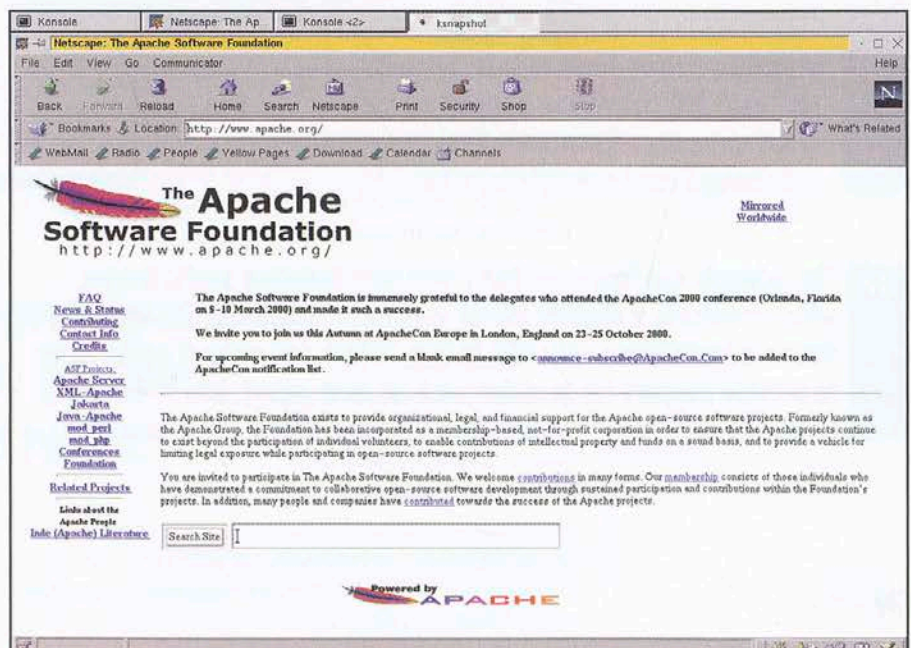
→ on bandwidth for the *Apache* folks. Anyway, download the latest version, which will have a filename such as "apache_1.3.12.tar.gz" into your /usr/src director, and unzip it with the following commands:

```
cd /usr/src
tar xvzf apache_1.3.12.tar.gz
```

Now you're ready to compile it all.

BUILDING THE SERVER

The first stage is to decide where we're going to put the *Apache* files themselves. While more traditional users prefer to install it in the original NCSA location of /var/lib/httpd, most users nowadays are opting for /usr/local/apache as the definitive location. We'll use the latter, although there's obviously nothing to stop you from installing it in, say, /export/apache or wherever – there are no hard-and-fast rules.

In order to compile, we need to set up the configuration for the relevant C-compiler and so forth. Thankfully, we can do this through a simple 'configure' script provided with the distribution, which provides us with all sorts of facilities to enable things like dynamic modules. At the moment, all we're bothered about is telling it where to install Apache, so change into your newly-created source directory and execute the script:

```
cd apache_1.3.12
./configure --
prefix=/usr/local/apache
```

While that's configuring, you'll see information scroll up the screen as the script generates the relevant makefiles. If all goes well, you'll see output like this:

```
# ./configure --prefix=/usr/local/apache
Configuring for Apache, Version 1.3.12
+ using installation path layout: Apache
(config.layout)
Creating Makefile
Creating Configuration.apaci in src
Creating Makefile in src
+ configured for Linux platform
+ setting C compiler to gcc
+ setting C pre-processor to gcc -E
+ checking for system header files
```

```
+ adding selected modules
+ checking size of various data types
+ doing sanity check on compiler and options
Creating Makefile in src/support
Creating Makefile in src/regex
Creating Makefile in src/os/unix
```

```
Creating Makefile in src/ap
Creating Makefile in src/main
Creating Makefile in src/lib/expat-lite
Creating Makefile in src/modules/standard
#
```

Now you're ready to build the server itself. There are no complex commands involved here. Just type:

```
make
```

And go and make yourself a cup of coffee while it's compiling...

Assuming no errors, you now have a compiled copy of *Apache*. Now there's ust one thing left to do, and that's to install *Apache* in /usr/local/apache:

```
make install
```

Finally, we'll test it by starting the server and checking that it works via our web browser:

```
/usr/local/apache/bin/apachectl
start
/usr/local/apache/bin/apachectl
start: httpd started
```
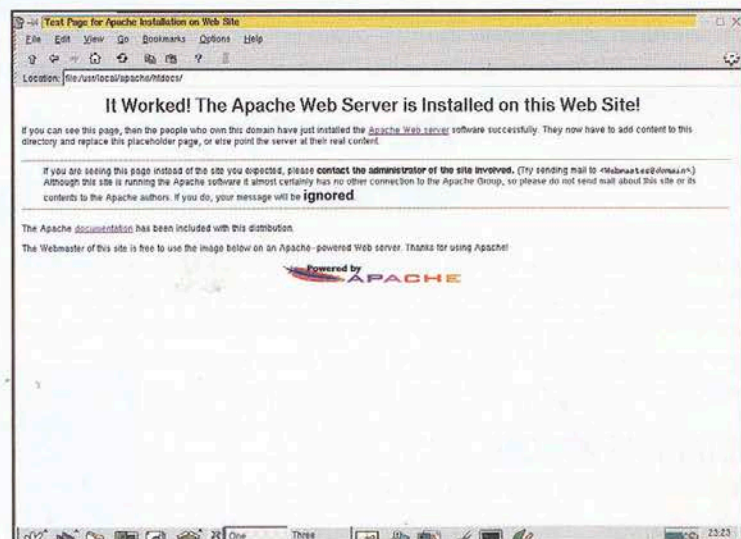

Hurrah! *Your* Apache *web server is set up and working if you see this screen.*

All server start/stop/restart operations are carried out via the "apachectl" script. It's worth putting this in your server startup script in order to start *Apache* when you boot. Assuming all is well, use the web browser on your Linux server to go to **http://localhost** and you should see the page shown above.

INSIDE THE APACHE DIRECTORY

Before we move onto configuration of the server, we'll take a look at the various directories which have just been installed in /usr/local/apache as some of these are of importance to us:
1. Program files are installed in bin/, include/, and libexec/. Manual pages are in man/ – Server configuration files are

installed in conf/ – CGI scripts are by default served from cgi-bin/.

2. Web pages are by default served from htdocs/ – the page you saw earlier when you tested the server is in this directory.

3. Icons for directory pages are stored in icons/ – Server logs are stored in logs/ – if you have server problems the first place to look is the error log in this directory, which can yield useful information. A log of page accesses is also kept here.

We don't need to worry about proxy/ for the moment, since we're not going to be running a proxy for a while yet!

CONFIGURING APACHE

While just putting your pages in /usr/local/apache/htdocs would provide you with a working webserver, it's bad practice and probably won't even work outside your own computer. Thus we need to configure *Apache* in order to ensure that we get the best from the software, and offer a degree of server security.

Since version 1.3 of *Apache*, the configuration has been centralised in one large file, httpd.conf, which lives in the conf/ directory. As with all sensible configuration files, it's in a textual format; load this into your favourite text editor and take a look. It may seem a little daunting, but it's very well documented and structured, having evolved over time to become flexible. For the moment, we'll just concentrate on the essentials of getting the server up and running. The file is split into three sections:

1. Global Environment
2. 'Main' server configuration
3. Virtual Hosts

For now we're more bothered about getting the main server up to scratch, so page down to the top of section 2 and we'll leave sections 1 and 3 for another time.

First of all, we need to change the "ServerAdmin" directive. This is the email address which appears on all error documents (such as the ubiquitous '404: Page Not Found' error). Change it to your own email address, or at least an alias:

ServerAdmin webmaster@jml.net

Now we need to make sure that *Apache* knows what the name of the server is, since in all likelihood it'll be different to the name of the server. This is important as it is used to create relative URLs: ServerName www.jml.net Note that you can't just invent names – they must resolve within the DNS. If you don't have a name which resolves, you should enter the IP address of the server here, like so:

ServerName 192.168.1.1

As we stated earlier, the server looks into htdocs/ when it is asked to serve a web page. If you're unhappy with this (and good practice suggests you separate data from programs) you
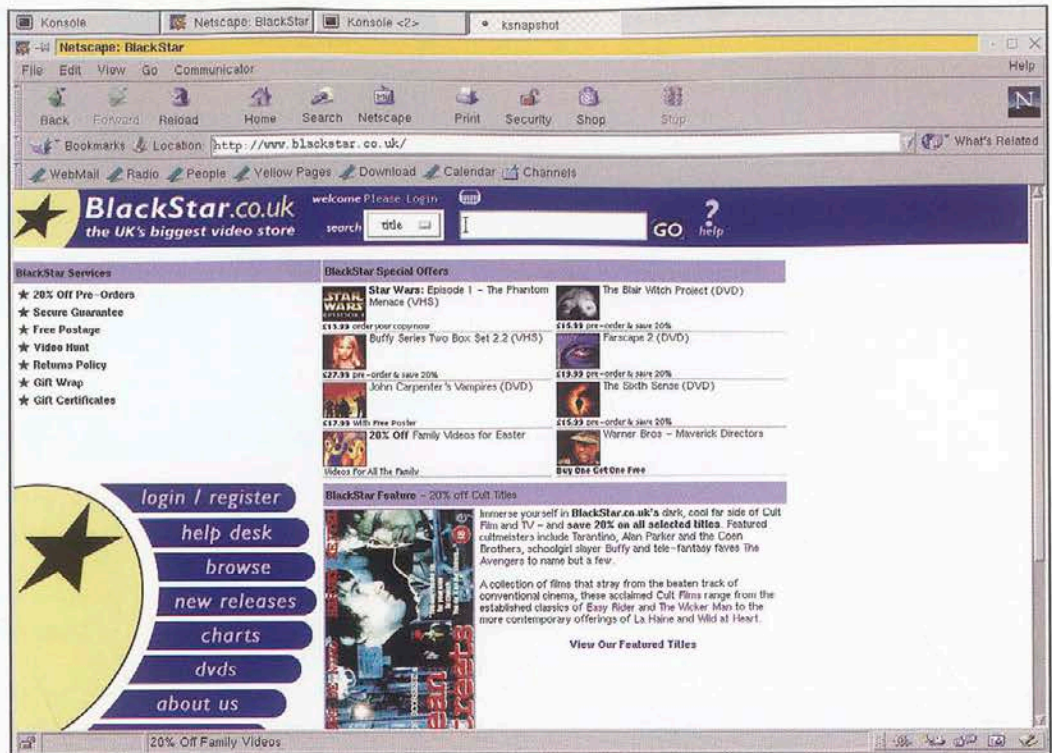
can change the location of the document root by altering the DocumentRoot directive accordingly:

DocumentRoot /export/webroot

Moving further down the configuration file to security settings, you'll find a series of directives structured in the following manner (this format is used throughout the *Apache* config and is referred to as a "container"):

```
<Directory "/export/webroot">
Options Indexes FollowSymLinks
AllowOverride None
Order allow,deny
Allow from all
</Directory>
```

The <Directory> container, as the name suggests, controls server permissions by directory. You will need to set at least one container for the DocumentRoot.

Although there are many different directory options which may be given, the two which have been specified on the "Options" line here are: Indexes: If no index page is found, *Apache* will automatically generate an index page as a list of files. If this directive is omitted, a "'403 Forbidden" error is given instead. FollowSymLinks: Symbolics links to other directories will be followed. Additionally, the "Allow from all" setting will not restrict anyone from accessing your website.

For the moment, ensure that the "/usr/local/apache/htdocs" entry is changed to whatever you set as DocumentRoot, ➔

Apache is very versatile, and forms the basis of many commercial websites.

You'll find that most source packages are distributed as GZipped TAR files. "tar xvzf {filename}" will unzip these (eXtract Verbally from Zipped File). TIP

We need to configure Apache to ensure that we get the best from the software.

→ otherwise pages may not be served.

Page down to the DirectoryIndex directive. This features a list of files which may be supplied as the default page for a directory. For example, if you go to **http://www.jml.net/** the actual file supplied will be **http://www.jml.net /index.html**. Multiple entries can be supplied on one line:

DirectoryIndex index.html index.htm index.shtml welcome.html

This example will check index.html first. If it's absent, a check will be made for index.htm and so forth. If none of these files exists in the directory in question then a "403 Forbidden" error will be given, or alternatively a directory listing subject to the "Options" line in the "Directory" container (as we saw earlier).

Finally, we're going to want to be able to run CGI scripts on our server from /export/cgi. We'd like /export/cgi to act as http://server.name/cgi-bin/, so we'll need to add in an alias of some kind. *Apache* makes this nice and easy with a directive called "ScriptAlias"

ScriptAlias /cgi-bin/ "/export/cgi/

We have to add in the relevant <Directory> container as well:

<Directory "/export/cgi">

The Apache logo must be the most common graphic on the web!

Apache supports all sorts of useful add-ons, such as PHP, Perl and CGI.

```
AllowOverride None
Options None
Order allow, deny
Allow from all
</Directory>
```

The ScriptAlias directive tells us that this directory is purely for running CGI scripts which, coupled with the prohibitive settings for the /export/cgi directory ("Options None"), should provide us with the security we need.

We're finished with the configuration file now. When you modify any aspect of the configuration, you'll need to restart *Apache* before the server takes any notice.

In order to restart *Apache*, you will need to issue the following command:

```
/usr/local/apache/bin/apachectl
restart
```

That's it – we've now got a basic web server which will quite happily serve a single website to other users, from the directory /export/webroot.

Although this is a very basic implementation of a web server, *Apache* is capable of a lot more, including the ability to host literally thousands of websites on one server and return content based upon database transactions, via technologies like PHP and mod_perl. Now have a play with it. Remember, information is power.

LINKS

The Apache Project:
http://www.apache.org
Apache Week:
http://www.apacheweek.org
WWW HOWTO:
http://www.linuxdoc.org/HOWTO/WWW-
HOWTO.html Apache SSL:
http://www.apache-ssl.org/
Matt's CGI Script Archive:
http://www.scriptarchive.com/
Analog, a logfile reporting tool:
http://www.analog.cx/

ANSWERS

If you are really stuck, why not write in? Resident gurus Dave Coulson and Richard Drummond will answer even your most complicated problems!

Hello!

Yes, I'm afraid it's time for another ridiculous picture of me, but more importantly, it's time for our Q&A section, where you ask the tricky questions and we answer.

There is a pretty broad spectrum of questions here, most of which we received in emails and letters in response to the Linux Answers magazine which was produced last year – which, cunningly enough, is one of the reasons we've decided to call this section Linux Answers...

Please feel free to email or write in with your own particular queries. The most interesting ones will be printed here every issue. We probably won't be able to answer your questions personally, so please don't be offended if you don't get a reply. If you want to stand a better chance of having your question featured, read the submissions advice boxout!

No question is too hard for us to track down the answer, so do your worst! Answering your questions are a team of experts including David Coulson and Richard Drummond. Oh, and I answer the easy ones...

Nick Veitch
Editor

Mounting Zips

Q I have an Iomega Zip drive that I want Linux to recognise. It's an internal Zip drive, and I think it is one of the original ones. It is plain IDE and not ATAPI. Any clues?
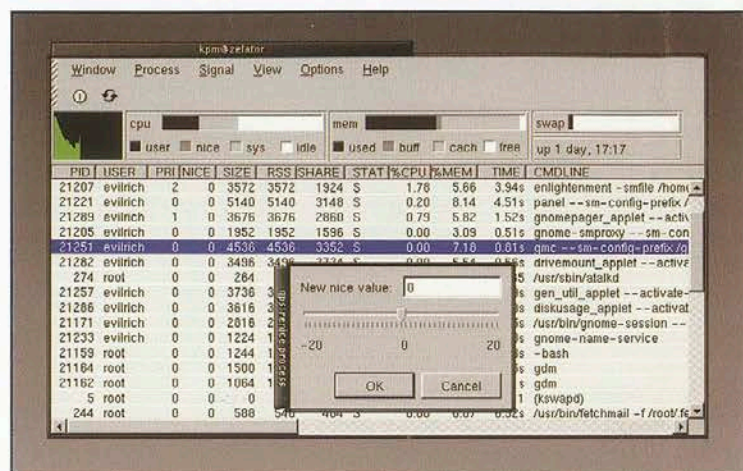
A The IDE Zip drive is handled in exactly the same way as a hard disk is in Linux, so you don't need ATAPI-Floppy support or anything like that in the kernel. You would initially think that all you have to do is 'mount /dev/hdd /mnt/zip', just like you would for a CD-ROM drive, but no, they have to be more awkward than that.

Most Zip drives actually have all the data on the fourth partition (though this might not always be the case). To mount the fourth partition on the disc at the mount point /mnt/zip, you would use 'mount /dev/hdd4 /mnt/zip', where /dev/hdd is the device the Zip drive is on. The directory /mnt/zip must exist for this to work, so if it doesn't, create it with 'mkdir /mnt/zip'.

Priority

Q Is it possible to prioritise tasks or applications running under Linux? I'd like to encode MP3s in the background while getting on with my usual work, but it seems to demand so much processor time that other things are painfully slow.

A Linux processes can have priorities ranging between -20 and 19 (-20 having the highest priority), but only the superuser can specify a negative priority or nice value. You can either initiate a command via the 'nice' program, setting a default nice level using 'nice -n ‹value› ‹command›', or if you'd rather change the priority of a current process, you can use 'renice ‹value› ‹pid›'.

So, in order to have your MP3 encoder running for less of the time, leaving more time for your other tasks, you want to give it a nicer value (that is, higher). Most tasks default to a nice level of 0 if you haven't changed the default. To find out the process number of your task, you can use the ps command. 'ps -A' will give

you a list of all the current processes and their pid.

A simpler way, if you are running KDE, is to use the KDE task manager. You can then select the task from a list and renice it with the option from the Process menu.

Make life easier on yourself: use a graphical tool such as kpm to manage your processes.

Installing

Q Where should I install stuff? What I mean is that whenever I download tar files and decompress them, and then compile source, I just end up with an executable in
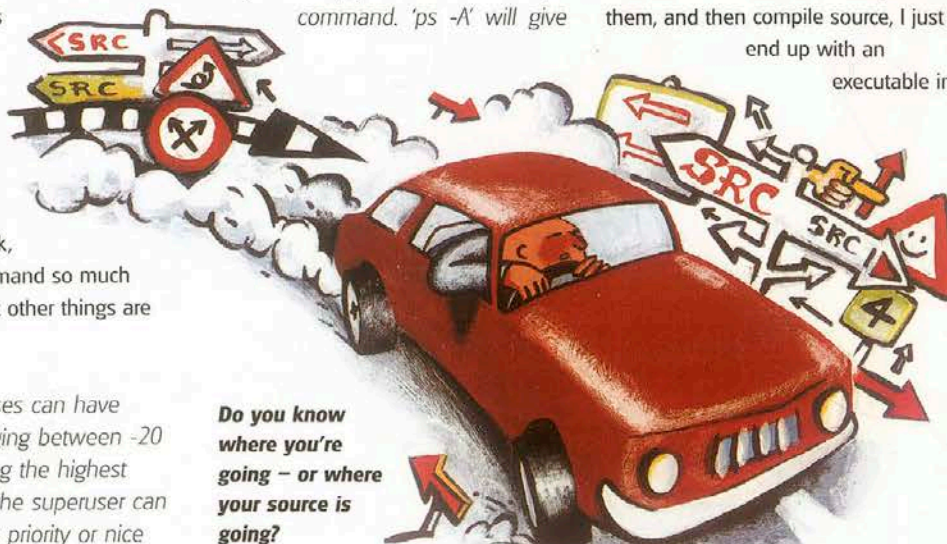
Do you know where you're going – or where your source is going?
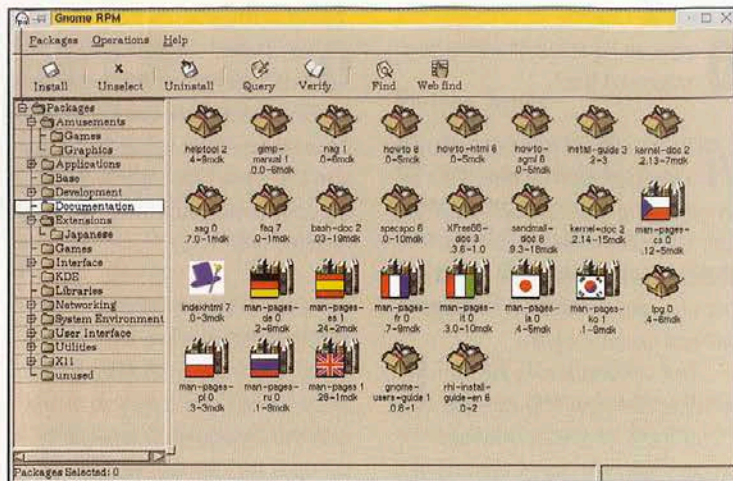
A package manager will save most of that mucking about with makefiles.

that directory. Where should it be, and whereabouts does Linux expect me to store the source? There seem to be loads of directories called "src".

A There are two general ways to install programs in Linux. The first is to use whatever package manager is supplied with your distribution. This will be RPM in the case of Red Hat and its derivatives, dpkg for Debian. This will provide you with a sophisticated level of management. It will install all the files belonging to a package in their correct places and it will also warn you if there any other programs that you are required to install first to make it work.

The second method is to compile the source code yourself and then install it. This is far more complicated than using a package manager, especially if you run into any problems.

Before doing anything else, make sure you thoroughly read any documentation supplied with the package. There will usually be files called README and INSTALL that will give you tips on how to compile and install it.

The standard way to build a source tree is to run './configure' from within the top level directory of the source, and then do 'make'. At this stage you'll end up with a tree with both the source and binaries,

which isn't actually much use. The final step, unless you want to do a 'make test', is to do 'make install', which installs the program, its libraries, headers and general files into it's default place, usually under /usr/local. Since most package managers have their program's install under /usr, if you don't remove the package-based version before installing from the source tree, you'll run into conflicts further down the road.

There is no set place to install source code, although most people tend to keep it in /usr/src, /usr/local/src, or simply in their home directory. As it's not needed to actually run programs, it doesn't really matter.

The exception to this rule is the kernel source code, which is needed to build a lot of programs. However, as long as you have /usr/include/linux linked to ‹prefix›/linux/include/linux, where the first linux is your main kernel tree, you won't have any problems.

Ghostscript

Q How do I use *Ghostscript* to display PostScript or EPS files? Is there a special version for PDF too?

A Ghostscript understands and will automatically recognise

PostScript, EPS and PDF files. Simply do 'gs filename' in X and it will open a window displaying that file. You would be better off using a graphical front-end for Ghostscript though, rather than just gs alone. Try GhostView or a similar package. For PDF files you can also use 'xpdf,' which

Don't use Ghostscript *directly*, use an X frontend such as AcroReader.

comes with most distributions, or which is available from this site:
http://www.freshmeat.net

Ethernet

Q I'm trying to set up my Linux box on the office LAN. It has a 3Com Vortex Ethernet card. Everything works fine under Windows, but I get a series of errors when trying to set up under Linux. It seemed to be recognised when installing Mandrake 7, but whenever I try to ping anything, or even when I try to connect to the Internet, I get errors telling me the "Network is unreachable".

I also get errors if I try to use the route command manually, telling me that the device is not ready, even though ifconfig is displaying the correct information.

A The Vortex Ethernet cards can be any one of a number of different models, all under the 3c59x name. The method to load the module is simply 'insmod 3c59x.o debug=1 option=X'.

You can set it to have a variety of different options, such as 10baseT, 100base-TX or 10Mbs AUI. A complete list is available in linux/Documentation/networking/vortex.txt.

If you still have problems, increase the debug level to see if it's having problems with any IRQ allocation. It may also be worth turning off Plug and Play in your BIOS, so the Ethernet card is allocated an IRQ at boot time by the BIOS and you don't have to mess around with the isapnp or kernel patches.

Archiving

Q I understand how to use tar and even how to untar a zipped file, but I notice there are quite a few files available that end with .bz2, which I understand is a completely different compression program.

With ordinary zipped files, I can just 'type tar xvzf xxx.gz' to uncompress them, but with the .bz2 files I have to uncompress them twice. Is there any way of doing this in just one line? →
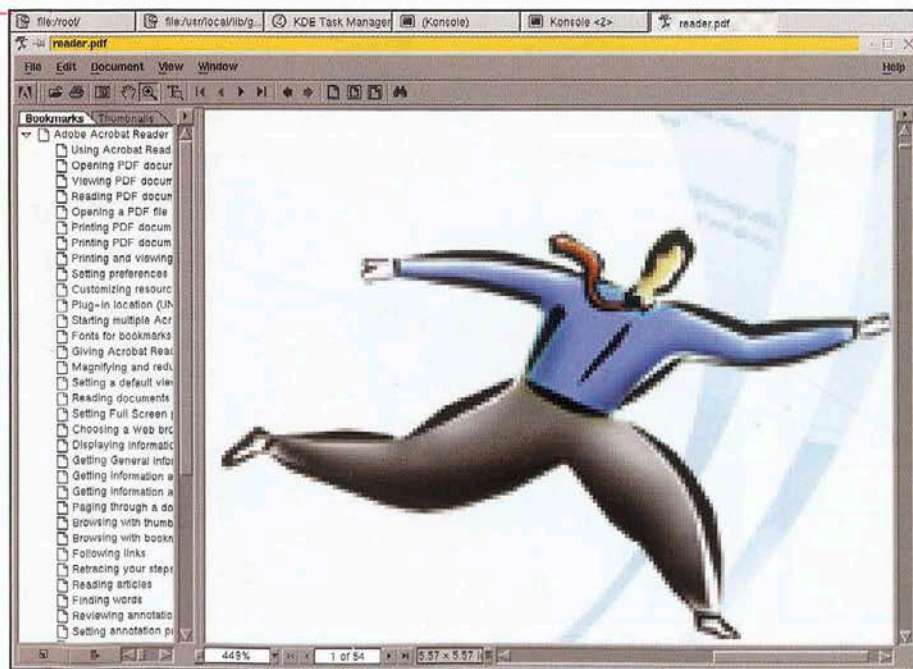
→ **A** The easiest way to do this is to get the bunzip2 command to unpack the file to stdout and use a pipe to pass the decompressed file to the tar command for unarchiving. Entering the line 'bunzip2 -c file.tar.bz2 | tar xvf -' in a shell should work.

XFree86

Q I have installed *XFree86* on my Linux machine, but I realise now that it is running in a particularly useless mode. What's the best way of reconfiguring it?

A The standard way to reconfigure XFree86 is via the XF86Setup command. First you'll need to have the machine in a text-based login mode, using 'telinit 3' (or 'telinit 2' if it's not based on RedHat) before you can run the command, as it will cause conflicts if you are already running an X server. XF86Setup will try to start the VGA16 server for its configuration window, so make sure that you have this installed.

Login

Q I have taken the advice of other Linux users I have met on the web and I now log on under my own user identity, rather than as root.

The only trouble is, now that I log on under my own identity, I can't seem to install software, edit configurations or do anything, really. How do I log on as myself, but have all the power of being root?

A Running as root is usually a bad mistake waiting to happen, since a misplaced asterisk or a wrong command could potentially wreck your whole installation. There are a couple of ways to achieve root privileges without actually logging in as root to do everything else.

The first method is to use the 'su' command, which will allow you to create a root session whenever you need to.

Another more powerful way is to use the 'sudo' command. sudo will let you set up various administrative privileges for particular users. You could let them run individual programs as root, or stop them running damaging programs as the superuser.

USB

Q My Linux distribution claims support for USB, but does this extend to printers? Also, how would I set it up? The printer doesn't appear as LP0.

A USB support in the 2.2 series Linux kernels is not particularly great unless you are using the backport of the 2.3 code. Within the new device structure the USB printers are under /dev/usb/lpX, where X is the number of the USB printer, starting at 0.

More information on USB in Linux (including the kernel patch) is available from this site: **http://www.linux-usb.org/**.

Using X

Q How do I quit out of *X* back to a command line?

A If you are not running a display manager such as xdm, then simply killing the X process itself will quit X. However, when using a display manager, killing X alone is not enough: the session manager will just restart it again.

The solution it is to kill the window manager. You can find the PID of the X or session manager process with the 'ps -A' command.

The other point to remember is that X runs in a virtual terminal and so all your other terminals should still be available. This means that you can use 'Ctrl-Alt-Fx', where x is a number between 1 and 6, to switch to a virtual terminal but keep X up and running. To get back to X, use 'Ctrl-Alt-F7'.

Cron

Q I understand that there is a program called cron which is used for running tasks at a specific time. I can't seem to find any information on how to set this up though. Can you help?

A There is a cron howto at the following website: **http://linux.com/howto**. You'll also find it at your local LDP mirror. If you'd prefer an offline source, 'man 5 crontab' gives you the syntax, along with numerous examples, of the /etc/crontab file.

Internet

Q My Internet connection isn't particularly great, mainly because my ISP is often slow. I was wondering if I could save time by setting up my own DNS server for pages I visit regularly. Could you tell me how I could do this?

A Depending on your distribution, there is usually a package called 'caching-nameserver' which installs bind/named for you, along with the configuration required to set up a DNS server locally. However, you'll still have to create an up-to-date list of the main root level servers, simply by typing 'dig › root.hints', then copying it to the right place.

Whether or not you'll actually get an increase depends on your surfing habits. If you return to the same site many times a day then it's a huge advantage, but for general surfing from one site to another you'd probably notice a fairly major performance drop.

Also, you can't save DNS requests between reboots, so every time you reboot, or even restart the name server, you'll have to request the DNS settings from the main name servers.

Backing up

Q Could you please tell me what the best way of backing up my Linux partition is? I have a CD-RW drive, but most of my Linux stuff is on a 2Gb partition.

Session   Language   Iconify

Welcome to zelator

G·N·O·M·E

Login

evilrich

Please enter your login

Remember that logging in as root can be extremely dangerous. It's much safer to only create a superuser session when you need it.
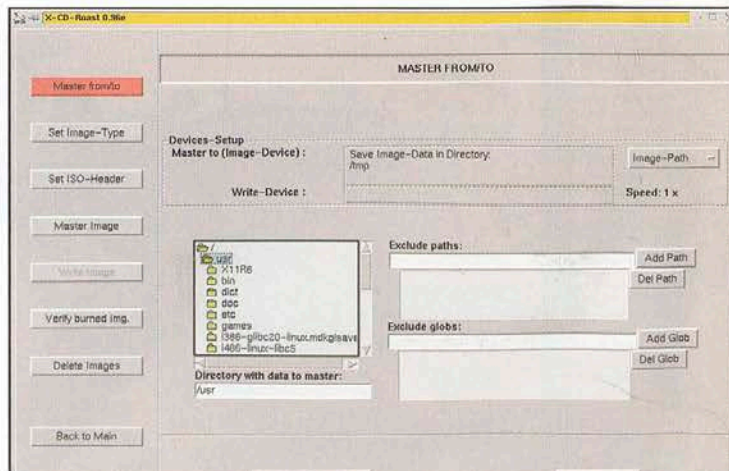
There are many CD burning applications for Linux, though not all are ideal for archiving tasks. Pick a backup method that suits your drive size and usage.

A There are many different tools available for backing up under Linux, so it really depends on your personal tastes and the medium you wish to back up to.

For burning the actual CD, cdrecord is ideal. You can either create a cd image on your hard disc and use rsync (http://rsync.samba.org/) to update it with various parts of your hard disk, then burn it to CD, or you can use a proper backup tool, including one of the many which are available at http://freshmeat.net/appindex/console/backup.html. Something along the lines of cdbackup (http://www.cableone.net/ccondit/cdbackup/) may be what you need.

Data Transfer

Q I have a lot of image files, text, etc, on a Windows machine that I developed for a website. I'd now like to transfer this to Linux so I can test the website properly under *Apache*. What's the best way to move the files without upsetting the file structure, filnames, etc?

A The easiest way, if the two OSes co-exist on the same machine, is to simply mount the Windows drive under Linux. This is pretty straightforward. You just need to find out the device name and partition number of the Windows partition, and mount it as you would any other drive, but with a partition type of

'vfat' (for Win 98 drives) or 'fat' for older partitions.

Failing fonts

Q My Linux system was working fine until recently, but then *xfs* started failing on startup and I'm having terrible trouble with *KDE*. Everything else works fine, but I can't see why *xfs* is failing.

A It could be a number of different things, but the usual reason is that you have been messing around as root and somehow changed the permissions for the /tmp directory. xfs needs to create a couple of directories in there for font management.

How much free?

Q I have several drives attached to my Linux setup and one or another of them is always running out of space. I know that I have a lot of stuff, and I have used *du* to find out how much stuff there is in various directories. However, what I need to know is how much free space is left on each partition.

A This one isn't so tough: you simply have to use du's relative, df (short for disk free). This will show you the percentage and amount of space free which is available on each drive, and much more besides should you care to know!
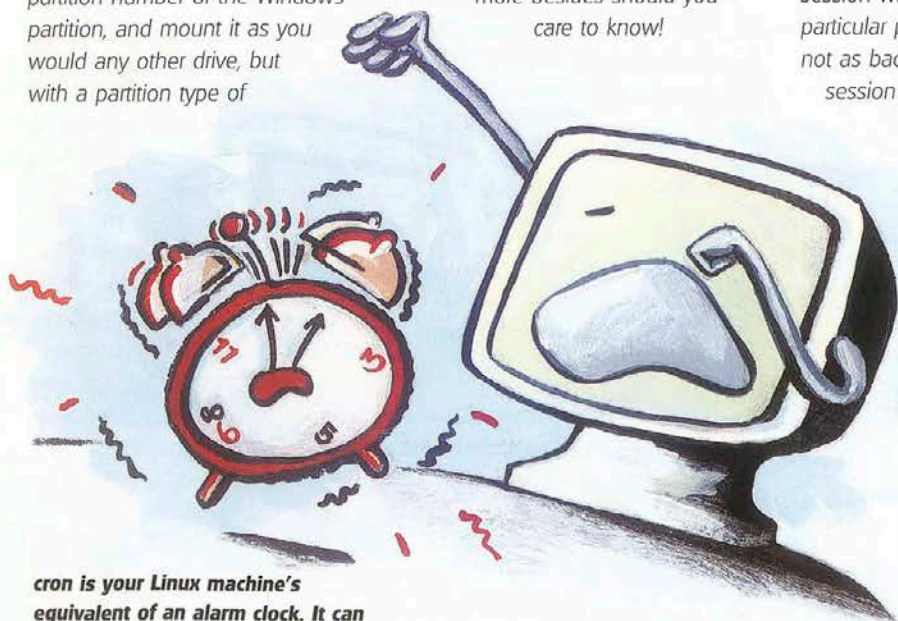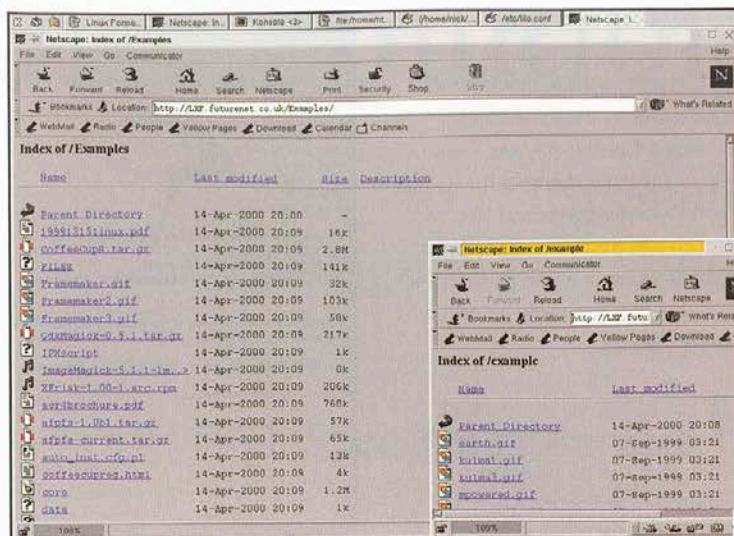
You can define your own filetypes and your own icons for Apache directory indexing.

Startup

Q How can I make a particular program run on starting up?

A It depends where exactly you want it to start. If you mean start up, as in boot time, simply add the command to /etc/rc.d/rc.local, not forgetting to add an & to the command to make sure it runs in the background.

If you want it started when you log in under X, add a line to run the command either to the file ˜/.xinitrc (if you launch X with startx) or the file ˜/.xsession (if you are using a display manager such as xdm, gdm or kdm).

Again, if you forget the &, the session will lock up until you kill the particular program. Fortunately, it's not as bad locking up your X session as it is locking up the whole machine during boot up.

Checking the mail

Q I use fetchmail to download mail from my ISP's mail server. Can I make it run automatically when I am online?

A The pppd (Point-to-Point Protocol daemon) is the system which allows

your Linux box to communicate with your ISP. When pppd successfully makes an Internet connection, it will execute the script /etc/ppp/ip-up, and when it closes a connection, it will then execute the script /etc/ppp/ip-down.

The secret to getting programs to automatically start up and stop as you go online and offline lies with these two scripts.

You shouldn't actually modify these scripts themselves: they should provide some method, depending on which Linux distribution you have, of running additional scripts.

With Debian, for example, the ip-up script executes everything that's in the directory /etc/ppp/ip-up.d/ and ip-down everything in /etc/ppp/ip-down.d/. So, in this case, you would add a script to start and stop fetchmail in the respective directory.

To start fetchmail as a daemon you would call it with a command like 'fetchmail -d 300'. The 300 is an interval in seconds to wait between checking its list of servers.

To stop fetchmail, you just have to use the command 'fetchmail -q'. Remember that if you do call fetchmail from the ip-up script then it will be run as root. Your fetchmail configuration file, .fetchmailrc, must then live in /root/. →

cron is your Linux machine's equivalent of an alarm clock. It can give any task a wake-up call.

→ Apache Images

Q I'm using *Apache* to run an Intranet. Actually, that is rather a grand name for three machines connected together. Anyway, I was wondering about changing the default graphics for directories and files.

Can I do this? I also have some filetypes which I'd like to give specific icons, but they just come up with a generic file icon.

A You can do all these things because Apache is very flexible, although it can be hard for newcomers to find their way around. The icons for Apache fancyindexing will normally be found in the /httpd directory, wherever that resides on your Linux system (usually in /home).

The default icons are all in the /icons drawer, and you can change these if you like.

The answer to the second part of your question is to edit the httpd.conf file, which is the general server configuration file for Apache. This will normally be found in the /etc/httpd directory (it may be called srm.conf, and be in the Apache directory on older Apache versions).

About a third of the way into the file, you'll find a lot of lines beginning like this:
Addicon /icons/binary.gif .exe .bin
Addicon /icons/binhex.gif...

The trick is just to add your

Your Linux box can easily share a printer with NetWare machines.

own line here. If, for example, you want to provide a specific icon for PDF files, you would add:
Addicon /icons/pdf.gif .pdf

This would mean that all files ending in ".pdf" would be given the icon pdf.gif, which is found in the httpd/icons drawer.

You must remember to put the icon in the correct drawer, and make sure the permissions are set to allow everyone to read it.

You could also edit these lines to point to your alternative icon set, but again, you have to make sure that the permissions are set correctly and that the other icons reside in a path where the server allows access.

Don't create command aliases every time you log in, just edit your .bashrc file and they will be available for every shell process.

Vanishing Alias

Q I have got the hang of using the Alias command and I've set up some very useful ones, but of course, every time I restart Linux I have to set them up again. I have thought of adding them to the init.d files but I guess that would be wrong, and also unlikely to work, because that's run as root isn't it? Any words of advice here?

A You are right that you shouldn't put them in with your init scripts – the correct place for this is the bash configuration file (assuming you are using Bash, that is), the .bashrc file.

You may not have noticed this file because it is usually hidden by the system (it begins with a '.') but it lurks in the home directory of every user.

The .bashrc file is just a script and it is run every time a bash shell is generated. This is important, because it means that anything
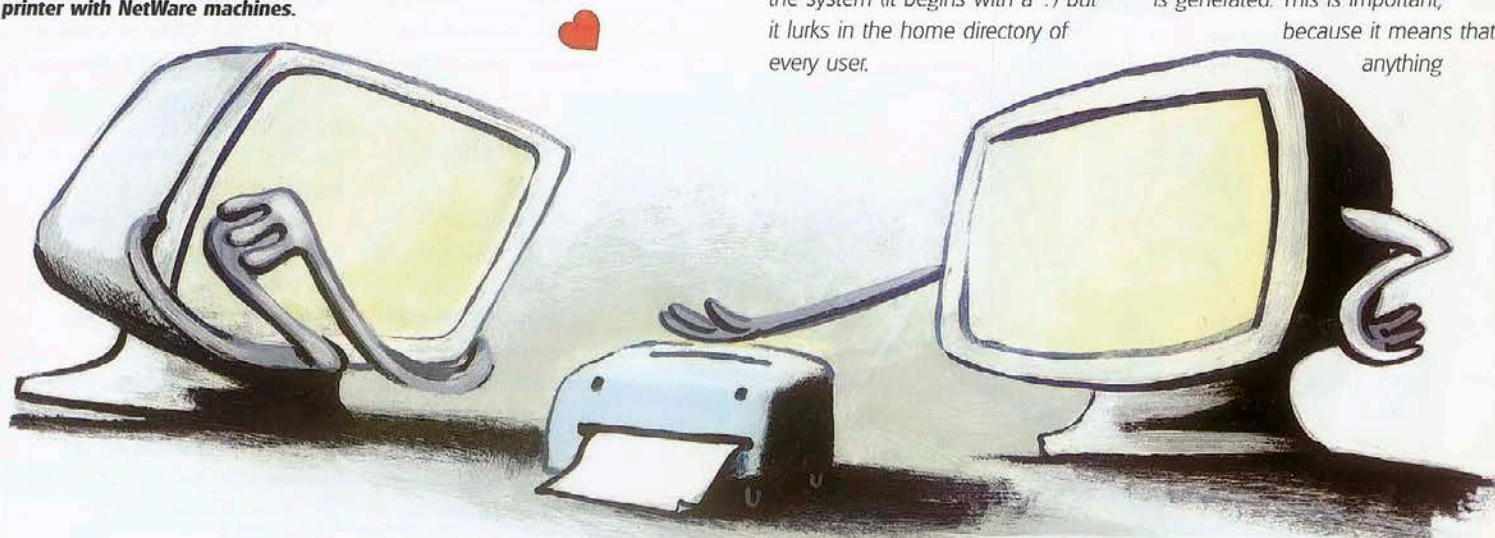
specified here applies to every shell that is opened, including ones that are spawned as the result of you running scripts.

As the .bashrc file is just a script, you can edit it to do anything you like. There are some commands in there which are best left alone, but you can tag your aliases on with the others which will undoubtedly already be there.

There is also another, slightly less useful file associated with a bash session – the .bash_logout. This file is not always automatically created for you, so you may have to just generate it with a text editor and save it in the same place as the .bashrc file.

As you might imagine from the name, this script is executed whenever you logout from a bash shell. It can at least be used to clear the screen, or perform any specific cleanup tasks before the next user logs in.

Networking

Q We are running Linux on a LAN, at the moment, in conjunction with NetWare. Is there any way in which we can share printers between the two networks?

A Of course. NetWare machines communicate via the IPX (Internetwork Packet Exchange) protocol rather than IP, so the first step is to recompile your kernel with support for IPX. To able to do anything useful over an IPX network, you will also need to add ncpfs support to the kernel. This will provide the means by which your Linux box can access services on a NetWare machine and vice-versa.

You will then need to install the ncpfs package on your machine. This package contains various client and server programs and which ones you need will depend on whether the Linux machine will be acting as a print server or as a workstation printing to the NetWare network. It should be available as a part of the Linux distribution you are using. For example, the ncpfs client, nprint, will allow you to insert files

into a remote printer queue on a NetWare print server.

A complete overview of how to set all this up is available from **http://linux.com/howto/IPX-HOWTO.html.** We will also be covering this in greater depth in next month's tutorial section.

SUBSCRIBE

DON'T DELAY – SUBSCRIBE NOW!

Title _____ Initials _____ Surname _____

Address _____

_____ Postcode _____ Country _____

Email Address _____ Daytime telephone number _____

☐ Please start my subscription at the rate ticked below

Please choose your method of payment **1, 2, 3** *or* **4**

1. ☐ Direct Debit (UK only) – I wish to pay £24.95 every 6 months.

I understand that I will receive 13 issues every 12 months.

ORDER Code No. LNXP02

Instruction to your Bank or Building Society to pay Direct Debits.

future PUBLISHING   **DIRECT Debit**

Future Publishing Ltd. Cary Court,
Somerton. Somerset. TA11 6BR

Originator's Identification Number

9 3 0 7 6 3

Please send this form to address shown

1. Name and full postal address of your Bank or Building Society branch

To: The Manager

_____ Bank/Building Society

Address _____

_____ Postcode _____

2. Name(s) of account holder(s)

3. Branch sort code (from the top right-hand corner of your cheque)

4. Bank or Building Society account number

5. Instruction to your Bank or Building Society

Please pay Future Publishing Direct Debits from the account detailed on this Instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s) _____ Date _____

Ref No (Office use only) _____

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

LINUX FORMAT

The UK's ONLY dedicated Linux magazine

Cheque / credit card	13 issues
UK -	£49.90
Europe (incl. Eire) -	£73.53
North America -	£73.53
Rest of World -	£86.53

2. ☐ Cheque (payable to Future Publishing Ltd. Sterling cheques drawn on a UK account.)

3. ☐ Visa   ☐ MasterCard

4. ☐ Switch   Issue No/Valid date _____

Card No. _____ Expires _____

Signature _____ Date _____

UK readers return this coupon to:
LINUX Format, Subscriptions, Future Publishing Ltd,
Freepost BS4900, Somerton, Somerset, TA11 6BR

Overseas readers return (postage payable) to:
LINUX Format, Subscriptions, Future Publishing Ltd,
Cary Court, Somerton, Somerset, TA11 6TB, UK

Internet address: http://www.futurenet.com/promotion/pg002/92

☐ Please tick this box if you would prefer not to receive information on other offers.

ORDER HOTLINE ☎ 01458 271178

QUOTE ORDER No. LNXP02

for only £24.95*

Pay £3.83 instead of £4.99 an issue when you subscribe to LINUX Format.

You have a choice: you can trudge down to the newsagents every four weeks and pay £4.99 for **LINUX Format** (great value in itself) – assuming, of course, that the magazine hasn't already sold out.

Alternatively, you can sit back in your comfy chair and pay as little as £3.83 an issue which is delivered to your door for free (UK only). You'll never miss another issue, and you won't have to give yourself a hernia getting the magazine each month.

RISK-FREE GUARANTEE You can cancel your subscription at any time and receive a full refund on all unmailed issues. No questions asked.

BENEFITS: Save over **£1** off every issue*. Keep your money in your pocket longer when you subscribe by Direct Debit

It's easy! All the details you need to complete the form are on your cheque. Delivery direct to your door. Never miss an issue

All you need to do is fill in the form on the left, and post it in a letterbox. If you live in the UK you don't even need a stamp!
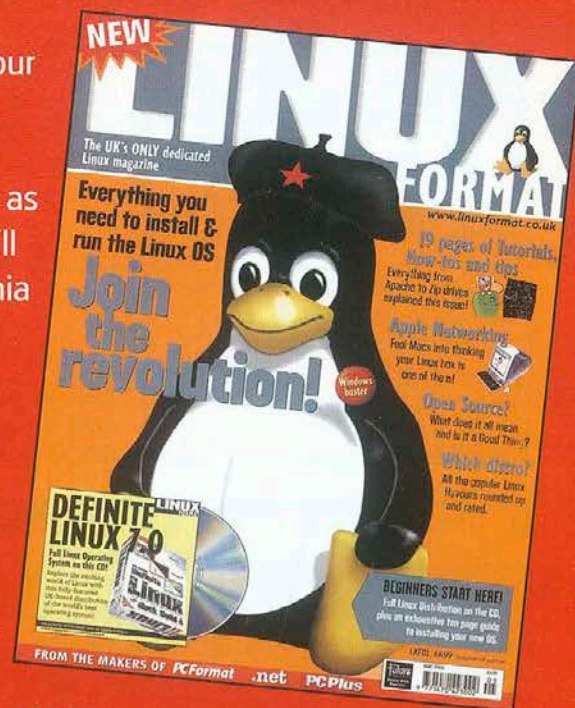
ORDER HOTLINE: Quote order code LNXP02 when you call   Lines open 24 hours a day

☎ **01458 271178**

* UK price only

LINUX Format Coverdisc

This month's CD features a full Linux distribution, **Definite Linux 7.0**, but please read the information below before you install

LINUX FORMAT on the CD

When you see this box in the mag, it means there's related stuff on our CD!

DEFINITE LINUX 7.0

Full Linux Operating System on this CD!

Explore the exciting world of Linux with this fully-featured UK-based distribution of the world's best operating system!

PRESENTED WITH ISSUE ONE OF LINUX FORMAT

For Details of supported hardware, please check documentation on CD.

Welcome to the very first Linux Format Coverdisc. It's a bit special because we have a full distribution of Definite Linux 7.0 for you to play with. Definite came out pretty well in our roundup of distributions and is certainly worth trying out.

It may use a fairly standard Red Hat installer, but I'm sure you'll find it no less easy to use - especially as we have a huge Newbie's guide this issue, which focuses around Definite 7.0. Even if you are installing Linux for the first time, don't be afraid! It is really a lot easier than you think, and you can check out our beginners guide starting on page 48

There is a guide in HTML format on the CD, with some useful links.

If you are installing Linux and you read and understand the installation guide, there's little to worry about. There is always some inherent risk in even turning a computer on, never mind installing a new operating system, so please don't just go running the installer and selecting things for the fun of it, as you may lose valuable data - don't come running to us if you reformat your Windows partition...

Of course, if you have a badly out of date version of Red Hat, you can take advantage of this distribution without having to install it. RPMS of all the packages contained in this distribution can be found in the Def/RPM/ directory of the CD, which can be installed in the usual way with your package manager, or through a shell with the command:

```
rpm -u xxxxxx.rpm
```

where the 'x's denote the package name.

You will also find that we have some (well, one page actually) of HTML on the CD. In future issues this will evolve into your self-contained guide to the CD, as well as providing links to important websites and other resources both on and off the CD. Well, it'll become clearer when we do it!
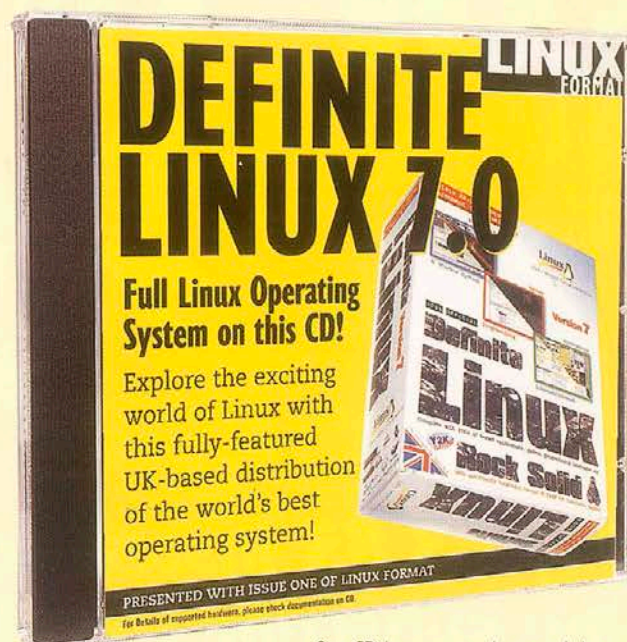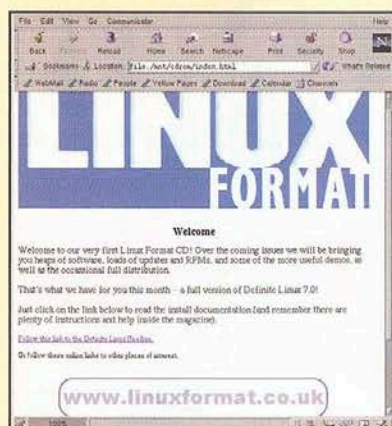
Our CD is very much a work in progress, and we need your input to make it better. See the boxout below on how to make it the best CD you could ever wish for – please take the time to fill out a survey! **LXF**

WHAT DO YOU WANT ON THE CD?

This is your chance to tell us what you want on the CD. Do you want updates, patches, latest software, HTML resources, PDF files, source, ridiculous pictures of the LXF team in embarrassing poses?

Whatever it is, and whatever we have to do to get it, we promise we'll try our best, but you have to tell us first!

You can do this in any number of ways. Firstly you could just write to us at **LINUX Format**, Future Publishing, 30 Monmouth Street, Bath BA1 2BW, or email us at linuxformat@futurenet.co.uk. The way we'd really like you to contact us though, is to use our on-line CD survey. Next time you pop by **www.linuxformat.co.uk**, just check out the on-line survey. You'll find the questions pretty easy to answer, and you'll be helping us to give you what you want!

IMPORTANT NOTICE

Before you even put the CD in your drive, please make sure you read, understand and agree to the following:

The Linux Format CD is thoroughly tested for all known viruses, and independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software before installing it. While every care is taken in the selection, testing and installation of CD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files before installing ANY new software.

JARGON BUSTERS

Not sure what we are going on about?
Do you get lost in the documentation that comes with your software?
Here's the solution – an explanation of the regular terms and software on the Linux platform.

Apache
The most popular Web server in the world, and it's included with most Linux packages. It's freely distributable and comes supplied with full source code.

Bash
The Bourne-Again Shell, used as the default shell on most Linux distributions. Bash features command and filename completion (press the <tab> key), and a searchable history of commands that have already been entered.

BIND
The Berkeley Internet Name Domain. A nameserver for Unix-like systems, which maps IP 0addresses to domain names.

Console
When you log in to Linux in text mode, you are at the console. Most machines run the X Window System, but on servers or older hardware which can't handle X, the console is the main Linux interface. You can have a number of virtual consoles, reached with the alt and Fn keys.

Daemon
This is a type of program which runs in the background, acting in response to certain events. For example, the Cron daemon can run programs every day or even every hour. Others, like Syslogd, keep track of various areas of activity on the system, and many end with 'd' (like 'pppd') to indicate the type of program.

Emacs
This is a very powerful text editor,

originally written by Richard Stallman of the Free Software Foundation. Extensions can be added using the LISP programming language, and other variants exist including *XEmacs*. It's popular among programmers, although some prefer *Vi*.

Ethernet
This is a very common networking technology, which can normally achieve speeds of up to 10 Mbps.

Ext2
This is the standard file system for Linux. Others, such as VFAT (used by MS Windows) and Amiga FFS are supported, although some of these file systems are still marked as being experimental. Ext2 features long filenames and extra code to reduce fragmentation.

Firewall
This is a method of security on networks, where all travelling data is scrutinised and is then either accepted or blocked, depending on a set of rules.

FTP
This is short for File Transfer Protocol and is used to send and receive files to and from a remote computer. There are numerous graphical clients available for Linux, such as *gFTP*.

GCC
The GNU C Compiler, a very powerful development tool which is commonly used for building kernels, etc.

GNOME
The *GNU Network Object Model*

Environment. Works with a compatible window manager to make a complete drag-n-drop desktop. The main feature is the Panel, which holds the application menu and can contain running programs. *GNOME* is the default desktop installed on Red Hat Linux, and it uses the Gtk library to create its user interface.

GNU
Stands for 'GNU's Not Unix', the operating system created by the Free Software Foundation. Linux systems today are comprised of GNU tools with Linux as a kernel, hence the name 'GNU/Linux', preferred by many.

GPL
The GNU General Public License, which requires that source code to a program is made freely available for modification. The Linux kernel is licensed under the GPL, as are most of the standard system tools and utilities.

Gzip
This is a compression program, similar to *WinZip* on MS Windows. It's typically used alongside the 'tar' program for creating archives.

Home
This is the directory where your personal files are stored. If you log in as fred, for example, your home directory will be /home/fred.

IP address
This is a unique number given to each computer on a network, such as '65.122.0.7'. This is the Internet Protocol address.

ISP
Stands for Internet Service Provider, a company like Freeserve which supplies a connection to the Internet for your computer.

KDE
K Desktop Environment. Built around the *Qt* libraries, this suite of tools includes a window manager, file browser and games. It's used as the default desktop for many distributions, including SuSE and Mandrake.

Kernel
This is the heart of the operating system, which communicates with hardware and shares resources between programs. The kernel for Linux systems was originally created by Linus Torvalds, and is now maintained by thousands of programmers all around the world. Torvalds still remains in charge of kernel development.

LILO
The *Linux LOader*. This small program sits on your hard disk and acts as a boot manager, allowing you to run Linux, Windows, DOS and other operating systems. *Loadlin* may also be used for booting from a DOS prompt.

Loopback
This is a network on your computer, allowing programs to communicate with it like a remote machine. It's required for running the X Window System.

Man
Manual page. Most command-line programs have a manual page, with info on usage and how to

contact the author, etc. For example, to get the manual page for the 'less' tool, you should enter 'man less'.

Module

This is a part of the kernel that can be loaded when required, rather than being built-in to the kernel image itself. This means that you get a considerably smaller and faster kernel.

Mount

This is the process of attaching a drive to a directory so that the files on it can be accessed. For instance, you usually mount the floppy drive device (/dev/fd0) onto the mount-point /mnt/floppy. Then you can look in the /mnt/floppy directory to see the files on the disk. A device should always be unmounted with the 'umount' command before being removed.

NFS

Stands for Network File System. This lets a user work with remote files as if they were local files, and NFS uses the TCP/IP networking protocol.

Open source

This is the name that is given to software which has its source code (ie, the original files in a programming language like C++) available for others to view. There is a range of open source licenses, the most popular of which are the GPL and BSD types.

Path

The directories that the shell searches when you enter a command. Usually this will include /usr/bin and /usr/X11R6/bin and can be seen by entering 'echo $PATH'. To run a program in the current directory, use './command' to show exactly where the program is located.

PPP

Point-to-Point Protocol. This is the most common method of making a dialup connection to the Internet, and an assortment of tools are available to make this process a lot easier, including Kppp and EZppp.

Prompt

Found at a command-line (in other words, in a terminal or logged-in at the console), and it is some text which reminds you of the directory you're in, which user you currently are and so forth. This prompt can be modified to include the date and time, among other details.

Qt

This is the name of a toolkit developed by Troll Tech which is used for building graphical applications. The KDE desktop environment is built around Qt, and it has come under fire from certain quarters because of licensing concerns (Qt isn't under the GPL).

RAID

Stands for the Redundant Array of Inexpensive Disks, where a number of small disk drives are used together, increasing performance and being able to tolerate problems better.

Root

This is the user who has access to all files on the system and is responsible for creating users, installing programs and other administrative chores. Also, the root directory is '/', which is where the main /usr and /etc directories start.

RPM

This is short for Red Hat Package Manager, which is a system for working with complete software packages. Files usually have a .rpm extension and can be installed using various tools, including KPackage and GNORPM. The system lets you check if other packages are needed before you can install new software.

Service

These are background processes; essentially another popular name for Daemon.

Shell

The command-line user interface. Various shells are available under Linux, the most popular of which is Bash. Shells create an environment for you to work in, and do this by providing shortcuts, scripting facilities and customised prompts.

Tar

This is a shortening of the term 'Tape Archive'. A Tape Archive is a program that joins together a series of files into one large file. Commonly used with the Gzip utility to compress the resulting file, it's the standard way of distributing source code. Such archives typically have a .tar.gz or .tgz extension.

Telnet

This is a program which enables you to work on a remote machine. With Telnet, you can log in and run programs over a network and display the results on your local system.

Tux

Tux has to be the most popular penguin on Earth, and he's the mascot for Linux. Tux was created by Larry Ewing.

UNIX

This is a multitasking, multi-user operating system which was developed by AT&T Labs in the late 1960s. There are many different flavours of Unix available, including BSD, AIX and, of course, Linux.

Vi

This is a powerful and initially difficult to use text editor, although it is preferred by many technical users because in bad situations it's the only one available. Popular variants include Elvis and ViM.

Window manager

This is a program for the X Window System, providing titlebars, icons and program menus for your desktop. Lots of window managers are available, ranging from the plain Twm through to the rather more extravagant Enlightenment. Others include IceWM, Window Maker and FVWM. The KDE environment comes with its own native window manager, which is called Kwm.

WINE

This term stands for 'WINE Is Not an Emulator'. WINE allows you to run MS Windows applications under Linux and other variants of Unix, although it's still in very heavy development and your mileage with it may vary quite a lot. Other alternatives include the VMWare virtual machine, DOSemu and Bochs.

XTerm

This provides a command-line within the X Window System. Other variants include Rxvt, Konsole and GNOME-terminal.

X Window System

These are programs and libraries that provide a graphical user interface for Unix-like systems. The version which is most frequently used on Linux PCs is XFree86, although commercial versions can be obtained. It's also known as X or X11.

YaST

Stands for 'Yet another Setup Tool'. This is the main system configuration program with SuSE Linux, enabling you to set up hardware and install software. It does a similar job to linuxconf on Red Hat-based distributions.
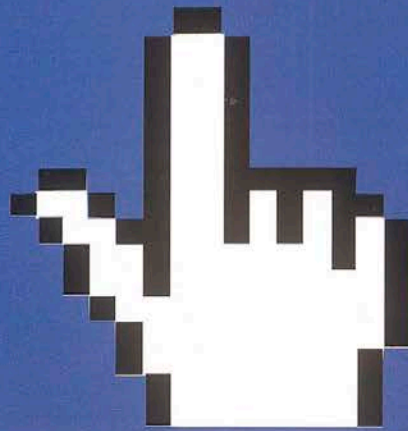
Zope

Stands for Z Object Pushing Environment. It's an open source web application server, used for building complex e-commerce sites with heaps of extra features.

Launches May 2000

NEW ECONOMY • NEW RULES • NEW LEADERS

BUSINESS 2.0

THE NEW ECONOMY STARTS HERE

Register for your free copy now
www.business2.co.uk

NEXT ISSUE

DON'T BE LEFT OUT! JOIN US NEXT TIME...

You've just reached the end of our first issue, but don't be despondent – there'll be another one waddling along to your newsagents in four weeks. And what's more, it's looking like it'll be even better than this one!

More features

Interesting Linux uses from around the world, the quest for a Linux PDA, customising your setup – these are just some of the topics we'll be covering in the next issue.

More reviews

We'll be taking a look at new boxed sets from Mandrake, SuSE and TurboLinux, the best new books, some interesting backup software and games. Also, don't forget our HotPicks section, which will tell you all about the latest and best software around!

More tutorials

Now you know a bit about writing shells, you can try your hand at shell scripting! We'll also have a look at more advanced *Apache* setup in our ongoing *Apache* tutorial, and there'll be more stuff to make and do.

More CD-type stuff

One thing's for definite, we'll have loads of software for you next issue, all categorised and easy to find, in various forms for maximum ease of use on your chosen system. Be here!

ON sale Wednesday 31st May

The contents of future issues are subject to change. We're not psychic, man, so don't get upset if some of this stuff doesn't happen!

THIS PENGUIN ROCKS

With the power of the AMD™ Athlon™ processor and the 3DFX Voodoo3 3000 AGP video card, there is no stopping this Penguin with attitude.

Our gaming desktop computer is equipped with today's hottest Linux games and home productivity applications.

Whether work or play, this desktop system is your high-end home Linux solution.

CRASH LESS – FRAG MORE
Superior gaming on Linux has arrived.

" If you crave a no-nonsense game station in the Linux environment, without having to be a guru, this is the machine you want. "

-Linux World

PENGUIN COMPUTING™

AMD Athlon PROCESSOR

The World's Most Reliable Linux Systems

Buy on-line or call 1-888-PENGUIN (888-736-4846)
www.penguincomputing.com
Tel: 415-358-2600   Fax: 415-358-2646

FREEDOM³

COREL® LINUX OS
Linux Power—Built for the Desktop™

www.corel.com/freedom