

EDITORIAL

Editor Nick Veitch
nick.veitch@futurenet.co.uk

Art Editor Julian Jefferson
julian.jefferson@futurenet.co.uk

Reviews Editor Richard Drummond
richard.drummond@futurenet.co.uk

Production Editor Andrew Channelle
andrew.channelle@futurenet.co.uk

CD Editor Paul Ravening
paul.ravening@futurenet.co.uk

Editorial Contributors

Rob Fenwick, Dave Coulson, Mike Saunders,
Steve Heaven, Jonathan Kent, Biagio Lucini,
Tom Wilkinson, Dave Cartwright, Alan Winfield,
Charlie Stross.

ART CONTRIBUTORS

Photography Rick Buettner, Louise Parker,
Kath Lane-Sims, Superstock Images Ltd
Illustration Gary Kempston, Garry Parsons, Tanya Myall

ADVERTISING

Group Advertising Manager Mike Roberts
Senior Advertising Manager Brian Ainge
Deputy Ad Manager Michelle Blackwell: 020 7317 2602
Display Sales execs Lorian Dorking : 020 7317 2630
Matt Dalton: 020 7317 2622
Classified Sales exec Tom Denning

ADVERTISING DESIGN

Point of contact/Team Leader David Mathews
Designers Dan Yeo, S Hobbs, M Stapleton, C Stenner
Supervisor Sarah Orchard

MARKETING AND PROMOTIONS

Marketing Manager Gill Stevenson
Marketing Exec Liz Britton
Subscriptions Manager Julie Sewell

PRODUCTION

Production Manager Lou Reffell
Production Co-ordinator Craig Broadbridge
Ad Designer David Mathews
Group Production Manager Judith Green

MANAGEMENT

Publisher Sheena Pittaway
Group Publisher Jessica Burley
Circulation Director Sue Hartley
Managing Director Mike Frey

DISTRIBUTION AND CIRCULATION

Circulation Manager Jamie Malley
jamie.malley@futurenet.co.uk

UK and overseas news-stand

Future Publishing Ltd 01225 442244

Overseas Licences

Tel: 01225 442244 Fax: 01225 732384

Contact Details

Linux Format, 30 Monmouth Street, Bath BA1 2BW
Phone: 01225 442244 Email: linuxformat@futurenet.co.uk

Subscriptions and Mail Order

Phone: 01458 271178. See page 84
Email: subs@futurenet.co.uk

Copyright No part of this publication may be reproduced without written permission from our publisher. We assume all letters sent – by email, fax or post – are for publication unless otherwise stated, and reserve the right to edit contributions. All contributions to *Linux Format* are submitted and accepted on the basis of a non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance, in writing. All CD-ROM demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary. We cannot be held responsible for mistakes or misprints. *Linux Format* recognises all copyrights in this issue. Where possible we have acknowledged the copyright holder. Please contact us if we have failed to credit copyright.

Disclaimer All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice. So there.

Printed in the UK by Midway Clark (Holt) and © Future Publishing Ltd 2001

LINUX is a trademark of Linus Torvalds.

Future Publishing Ltd. is part of The Future Network plc.

The Future Network serves the information needs of groups of people who share a passion. We aim to satisfy their passion by creating magazines and web sites that offer superb value for money, trustworthy information, multiple ways to save time and money, and are a pleasure to read or visit. This simple strategy has helped create one of the fastest-growing media companies in the world: today we publish more than 130 magazines, over 45 magazine web sites and eight specialist web networks from offices in seven countries. The company also licences 52 magazines in 26 other countries.

The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

Chairman: Chris Anderson
Chief Executive: Greg Ingham
Finance Director: Ian Linkins
Tel: +44 1225 442244

www.thefuturenetwork.plc.uk



Media with passion

BATH • LONDON • LOS ANGELES • MILAN • MUNICH
NEW YORK • PARIS • ROTTERDAM
SAN FRANCISCO • TOKYO • WROCLAW



WELCOME

Back in the days when *Linux Format*

was just a young nipper, an acquaintance of mine claimed that Linux was a bit of a nine-day wonder, and we'd soon run out of things to write. How wrong could they be? The Linux market is probably even more dynamic now than it was back then.

More than that, even in the UK Linux is being seen more and more as a seriously good OS, rather than just a hacker fad. The Open Source movement has more credibility thanks to high profile projects such *Mozilla*, which gave birth to *Netscape 6*.


The new Linux kernel 2.4 addresses many of the concerns over whether Linux is capable of scaling up to "Enterprise" level tasks, with greater support for symmetric multi-processing, and the addition of some key new architectures to the Linux family, the IA64 and S/390.

Desktop users have had a good year too. The proliferation of new easy-to-install Linux distributions, coupled with ever expanding support for common devices (such as USB printers) and, further use of Linux in PDAs (check out our preview of the Agenda VR3), makes it more attractive than ever. A flash in the pan? It would seem not.

This issue we are very pleased to be able to bring you the full source for the release version of the new kernel – 2.4. You should be able to happily upgrade your current kernel without fear, as this is no longer a "test" version, but the full stable release. If you are unsure, about how to do this, check out our instructions in the CD pages, starting on page 93.

If you are interested in what sort of things other people are using Linux for, look no further than page 14, where Alan Winfield tells us the tale of the Linux robots, and why the operating system is a crucial component for his work on behavioural robotics. And of course, there's plenty more exciting Hot Picks, reviews and tutorials for you to try in this issue!

Nick Veitch
EDITOR



See more of the *LinuxBot* on page 14.

WELCOME | intro

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 promote the use of Linux by providing friendly, easy to follow guides to installing and using this operating system.
- 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.

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

PCPlus

pcplus.co.uk

THE PC MAGAZINE FOR THE INTERNET



THE VERY LATEST PC TECHNOLOGY

HOTTEST HARDWARE

18 scorching upgrades to fire up your PC!

→ Graphics cards → Hard drives → Soundcards → Speakers

WINDOWS 2001

Microsoft Whistler revealed

MAXIMUM POWER!

9 killer PCs on test

MICROSOFT OFFICE 10

Exclusive full preview of Microsoft's new Office suite

MASTERCLASS

Top tutorials from our experts

- **Office** – step-by-step Access
- **Graphics** – photo-editing made easy
- **DVD** – create your own movie disc!
- **HelpDesk** – your questions answered



PCPlus new DVD and 3 X CD issues on sale **Wednesday 31 January**

PCPlus

INTERNET BUYER Your complete guide to buying online www.pcindex.co.uk




CONTENTS

Welcome to another jam-packed issue of **Linux Format**, your guide to all things Linux!

Walking Web

Enter the fascinating world of the Linux robots **p14**

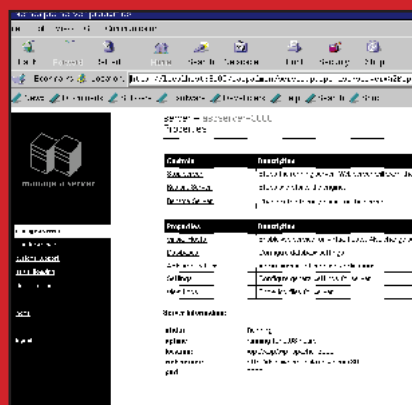
Alan Winfield introduces us to his Linux powered robots and tells us how and why they were developed using our favourite OS – and why Linux will continue to be used for the next generation of robots...



Chili!Soft ASP

Active Servers for Linux **p34**

The active server page market has quite a following, but funnily enough Microsoft's *Visual Basic* extension was never ported to Linux by them. Fortunately, Chili!soft's *Apache* extension makes it usable on Linux anyway...



Mastering Sawfish

This powerful window manager will let you reshape your desktop **p68**

Biagio Lucini reveals how a little librep knowledge can go a long way when it comes to giving Linux a makeover.



REGULARS

NEWS 8
All the news that's fit to print in our monthly roundup of the most important stories!

SUBS 84
Subscribe to **LINUX Format** and save money, never mind shoe leather. All the details are here.

COVERDISC PAGES 93
We've managed to squeeze an obscene amount of Hot Picks, kernel source and a whole distro onto this month's CD!

NEXT MONTH 98
Does exactly what it says on the tin.

LINUX USER GROUPS? 90

You're never alone with a LUG apparently, and there are plenty of them – we're sure to have found one near you, wherever you live...



REVIEWS

JBUILDER 24
Borland's Java development tool goes forth.

CHEEP LINUX MEGAPACK 26
The latest installment under the microscope.

COBALT RAQ 28
It looks cool, but how well does it work?

AGENDA VR-3 30
We take an exclusive look at this miniature marvel.

WIN4LIN 32
Is this a worthy contender for VMWare?

BOOK REVIEWS 36
Another selection of top reads.

GRAPHICS ROUNDUP 38
Reviews to stir your creative juices.

HOT PICKS 46
Including a brace of web development apps.

LETTERS 58
Your views and comments.

TUTORIALS

WHAT ON EARTH IS LDAP 62
The hows, whys and wherefores of the Lightweight Directory Access Protocol explained. Find out how it could make your life a little easier.

PROGRAMMING PERL 72
Charlie Stross tackles a real world problem using Perl's regular expressions to perform complex find/replace actions on text.

MORE JAVA MAGIC 76
The magic of multi-threaded applications is unravelled in this month's thrilling installment – find out how one program can do several things simultaneously (unlike any of the LXF staff).

LINUX VARIATIONS 79
Simon N Goodwin reveals some of the quirks, problems and solutions he's uncovered during his exploration of past *Linux Format* coverdiscs.


LINUX ANSWERS 86
Our experts have been wracking their brains to solve the pressing dilemmas of *Linux Format* readers.



Expand your HORIZONS

Get your 3D graphics hardware working to its full potential under Linux, with our guide to setting up hardware graphics acceleration!


p52



CONTENTS | issue 11


LINUX Format CD

**FULL
KERNEL
2.4 SOURCE
ON OUR CD
THIS ISSUE!**



PLUS...

Full TurboLinux 6 distro – find out why this is such a hit with sysadmins. Plus a selection of Hot Picks from this issue and more!



Please read the coverdisc instructions starting on page 93 before installing from the CD!



The Christmas lull is officially over, and all around the world developers and manufacturers are getting back to the serious business of releasing loads of new Linux hardware and software to fill our magazine. Well, they might be hoping to make some money as well, but whatever the reasons, if you want to find a way to wisely spend that £10 postal order that granny sent you, or your boss has given you £5000 to spend on a new web development box, we think you'll find the answers to your spending dilemmas within...
Richard Drummond

REVIEWS

All the very latest software, hardware, games and books put to the ultimate test – the scrutiny of our **LINUX Format** reviews team!

IN THIS ISSUE...

FUNKY COOL BLUE SERVER

COBALT RAQ4

The ISPs favourite webhosting appliance on test **p26**

Recent acquired by Sun, Cobalt have an outstanding reputation in the web-server appliance market. Join us to take a look at the workings of their flagship, the RaQ4.



THOSE OTHER REVIEWS

JBUILDER24

Borland's latest version of their Java development suite adds much needed features, but is it worth the cash?

CHEEP LINUX MEGAPACK26

The latest version of the Cheep MegaPack rounds up the most popular distros and delivers them to you in one bundle.

WIN4LIN32

Run Windows applications under Linux, but only if you manage to install the software! Get the lowdown on this new emulator.

CHILISOFT ASP34

Pesky VisualBASIC scripting additions, but running under Apache thanks to this commercial extension.

BOOKS34

The latest hot off the presses tomes of wisdom for your perusal.

GRAPHICS PACKAGE ROUNDUP38

The best the Linux world has to offer when it comes to Image manipulation and editing, rounded up and sorted out.

HOT PICKS46

LINUX Format delves into the open source community and comes up with another basketful of killer apps.

LINUX FORMAT



Our award for "TOP STUFF"

It's quite simple. If we really, really like something – I mean really think that a particular piece of software, hardware or any other sort of ware is the best stuff around – then we'll give it our "Top Stuff" award. Only the very best will be chosen.

HANDHELD HEAVEN?

AGENDA VR-3

As developer versions of the Agenda Linux based PDA ship, we take a sneak preview **p28**

The developer edition of the Agenda VR3 handheld has started to ship, so we got a PDA developer to give it the once over and let us know what the future holds for Linux handhelds.

JBUILDER 4

Richard Drummond checks out the latest release of this Java RAD.

ENTERPRISE EDITION

Publisher: Borland **Web:** www.borland.com **Price:** £1,999

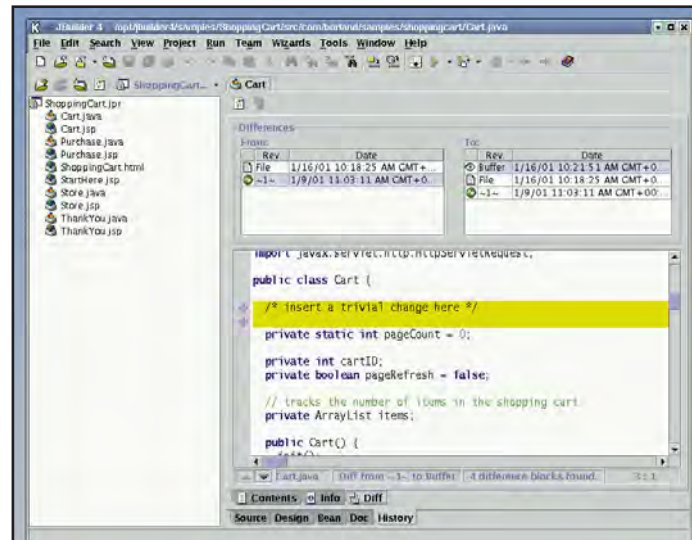
With the release of the Enterprise Edition of Java 2, Sun has yet again re-invented their Java platform, this time as a server-side application environment. And while Java might have failed to set the desktop alight, it is gaining rapid acceptance in enterprise roles because of the ease and speed with which it lets developers create distributed and web-based applications.

To take full of advantage of Java in enterprise, you need more tools than just a text editor and command-line compiler, though. Borland's *JBuilder 4, Enterprise Edition* is a complete visual suite that lets you rapidly build, debug and deploy enterprise-level Pure Java 2 applications. *JBuilder 4* is fully J2EE-compliant, supporting all the key J2SEE technologies such as Java Servlets, JavaServer Pages (JSP), Enterprise JavaBeans (EJB), XML and Java Database Connectivity (JDBC). Version 4.0 sees Borland consolidating their position as the leading Java RAD tool and improving support for EJBs, web and team-based development.

The basics

The core IDE in *JBuilder* is the AppBrowser, a multi-paned window where you can simultaneously manage the files in your project, edit source code, design GUIs, watch your applications run, view web pages and more. It will be instantly familiar to anyone who has used any other IDE, such as *Visual Basic* or *Delphi*, and has a similar functionality throughout the *JBuilder* range. If you want to try it out, then the Foundation edition of *JBuilder* is available for free download from the Borland website.

Personally, I don't particularly like the layout of *JBuilder*'s AppBrowser. For instance, with everything organised into groups of tabs, it makes it frustratingly impossible to view two source pages at a time without creating a whole new browser window. I would much prefer to be able to simply tear off tabs into floating windows whenever I felt like it. Moreover, I find the whole interface sluggish enough to be annoying even on a fast machine. Despite such niggles, it does have some great and new features.

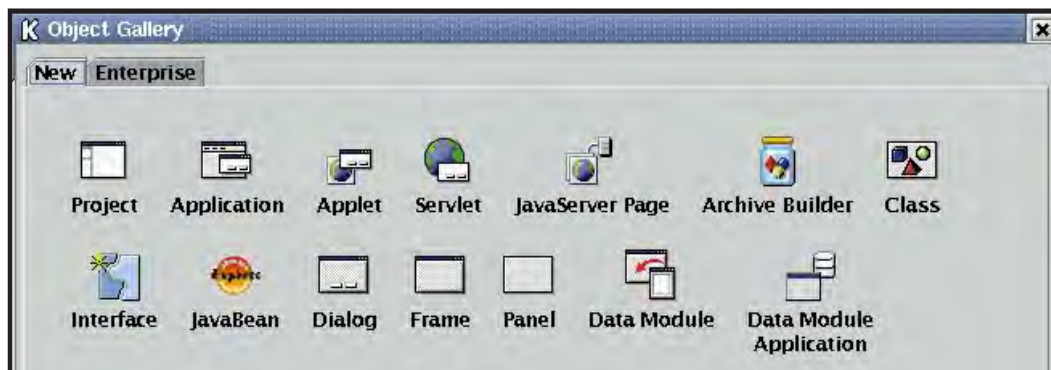


The new history view lets you examine the changes made between different code revisions.

The editor naturally enough supports syntax-highlighting, but the icing on the cake is the Code Insight tool. This pops up a menu as you type along and offers possible completions in the class library of the identifier you are currently typing. You can select the completion you need, or you are free to ignore it and keep on typing. Another nice touch is Structure Insight, which breaks your source code grammatically into a tree of elements and displays this in a

panel next to your code. It's great for quickly navigating through your code and provides a very visible check that your code is structured how you want it. In *JBuilder 4.0*, the syntax highlighting, Code Insight and Structure Insight have all been updated to handle XML, XSL, JSP and IDL source code.

The biggest new feature to the AppBrowser, however, is the addition of Team development. *JBuilder* now integrates with CVS to provide



Wizards can speed up the creation of most common objects.

SYSTEM REQUIREMENTS

JBuilder is coded entirely in Java and so the system demands are quite high. The minimum requirements are a 233MHz Pentium II or compatible, 128Mb RAM and 250Mb of disk space – although, you'll need to patient to run it on such a system. Realistically, 256 Mb RAM is required and the faster the processor the better.

concurrent version control of your projects. This allows multiple developers to work simultaneously on a project without treading on each other's toes and is a great boon for project management. You'll have to install and set up your CVS repository externally to *JBuilder*, but once that's done, you can use the GUI to create modules, check in changes, create diffs and view previous versions. This history view function will actually work without a CVS Server, too, instead using local backups of your source.

Full of Beans

A flashy IDE is all very well, but what about tools for Enterprise development? The fundamental technology of the J2EE is the Enterprise JavaBean, the server-side component for creating portable middleware. *JBuilder 4.0* support the upcoming EJB 2.0 specification and is chock full of functions to automate the creation, handling and deployment of EJBs. There are wizards to create session and entity beans, specify their interfaces and properties, and create test applications for debugging.

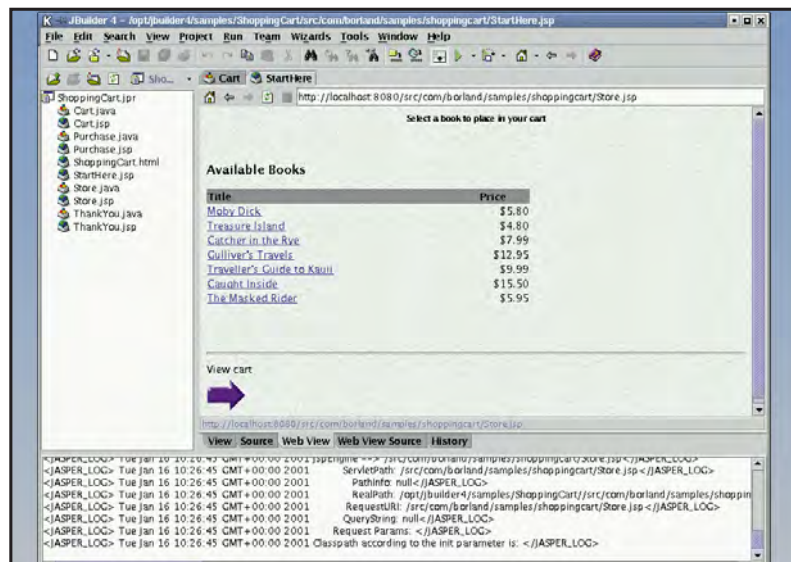
The default bean server shipped with *JBuilder* is the Inprise Application Server (although BEA's *WebLogic* server is also supported) and provides a means to deploy and test your beans – a developer only licence only is supplied. IAS is built upon Borland's *VisiBroker* for Java, so creating

CORBA applications is also fully supported. The supplied Caffeine compiler will automatically turn your Java code into IDL.

JBuilder features tight integration with IAS. For instance, it allows remote debugging of beans, and supports IAS's Deployment Descriptor Editor. The deployed descriptor is an XML file that describes how your beans will be packaged and deployment and is used by the EJB Deployment Wizard to automatically pack your Beans into jars.

Web development has also been enhanced in this version of *JBuilder* with support for the JSP 1.2 and Servlet 2.2 specifications. New wizards are included to automate the building of JSPs and Servlets classes. Also new is InternetBeans Express, a set of components to generate dynamic web content from template pages and DataExpress DataSets and DataModules. *JBuilder* will also generate HTML clients for CORBA applications.

This release of *JBuilder* sees the inclusion of the *Tomcat* 3.1 – Sun's reference implementation JSP and Servlet engine – which has been integrated into the IDE to the ease the testing and debugging of your applications. Running a web-based application will launch a new instance of *Tomcat* on an empty socket and serve up your pages. Alternatively, you may deploy your application to a full web server and test there.



The integrated web view displaying a web-based application.

INSTALLATION

The first thing you notice about *JBuilder 4.0* when compared with 3.5 – on Linux at least – is that it is much easier to install. Previously you had to download and install your own JDK first, because there wasn't one included in the package. With release 4.0, however, Borland have included IBM's JDK 1.3, which is set up to run from the CD and start the installation process.



Installation has been improved in this release.

Database development is taken care of by a development version of *JDataStore* – Borland's pure Java, embeddable database engine – and, for heavyweight applications, *InterBase 5.JDataStore* has received a number of enhancements with this release, including a significant improvement in speed and security. It can also now be accessed through JDBC or with Borland's DataExpress

components. The latter allows you to work offline, edit data and resolve the edits when next reconnected to the data source, and to store objects as well as data. Another new addition is the graphical browser, *JDataStore Explorer*, which has been supplied to complement the JDBC browser, and can be used to create and manipulate *JDataStore* tables. The package also includes Borland's *dbSwing* components, which are database-aware versions of the standard *Swing*

components. These improve the package's abilities in the rapid creation of database front-ends.

AYE, AYE, CAPTAIN!

JBuilder 4.0 Enterprise is a solid package and contains a number of excellent tools for building database-driven, distributed and web-based applications. It is a serious piece of software with a serious price tag, but will probably save you money in the long run by driving down development time. Of all the new, improved features in *JBuilder 4.0*, the ones that particularly impress me are the CVS integration and the new deployment wizards which should improve development time. **LXF**

LINUX FORMAT Verdict

Installation: 8/10
Ease of use: 8/10
Features: 9/10
Documentation: 9/10

Version 4.0 of this stalwart RAD tool sees productivity climb ever higher.

Rating 9/10

CHEEP LINUX MEGAPACK

Tom Wilkinson looks at a good way to get hold of the latest distros without tying up your phoneline for four days.

Publisher: Cheep Linux **Web:** www.cheeplinux.co.uk

Price: £ 29.99

Newcomers to Linux can often be confused by the fact that unlike many other operating systems, there is no single company who supply Linux. Frequently each company's distribution of the increasingly popular system is different from the others, in ways that vary from the subtle to the blindingly obvious. New users are well recommended to try as many different distributions as possible, and to settle for the one with which they are most comfortable. However, one problem that can arise here is that people may not know what is available, and hence make a bad choice. Cheep Linux, a distributor of a full range of Linux CDs for a low price, has come up with a possible solution to this problem in their *Megapack*.

The *Megapack*, which is revised regularly, contains a number of the

latest Linux distributions that have been selected on the basis that they are easy to install and use, as well as the *Definite Linux Bookshelf* CD. This contains a large number of manuals including the *KDE* and *GNOME* users guides, a guide to getting started, the HOWTOs and the *Linux User's Guide*, in addition to a number of more technical texts such as the *Linux Programmer's Guide*. To this, the pack also adds reader software for Windows, Mac and Linux so that they may be accessed before and after Linux has been installed.

Numbers game

At the time of writing, the pack contained five distributions: Mandrake 7.2, Red Hat 7.0, SuSE 7.0, Slackware 7.1 and Corel 1.2. This is a lower number than have been in previous editions of the package, because



Mandrake is one of the easiest to install distributions in the Megapack.

both Red Hat and Mandrake have recently swelled to two disks as opposed to the usual single CD. While this means that there is a slightly smaller choice, there is a greater variety of software available. Even if the distribution you finally settle on is not Mandrake or Red Hat, the additional programs, for the most part, may be installed as everything is supplied in RPM format. All the distributions supplied have scored above-average results in their reviews in previous issues of *Linux Format*.

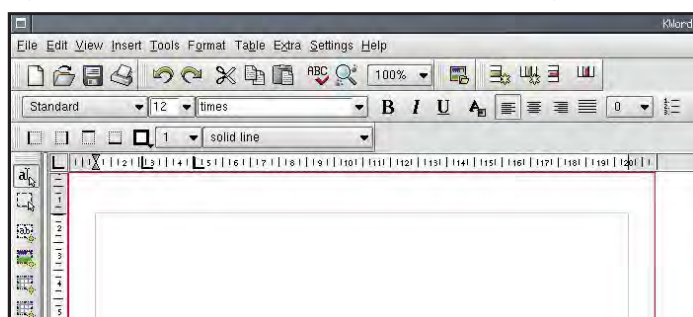
In addition to the eight CD-ROMs, the box includes a voucher which will let the user to obtain a further three CDs, paying only for postage. These can be selected from distributions which are not included in the box, such as Debian or

TurboLinux, or from disks full of applications including games and utilities. There is even a disk full of desktop themes. The voucher can also be used at a later date to obtain a new, up-to-date version of a distribution included in the box. All considered then, the *Megapack* offers a good value 'starter kit' for relative newcomers who are ready to experiment with a number of different distributions. **LXF**

LINUX FORMAT Verdict

If you're looking for a collection of distributions to help you find the one you're suited to, the *Megapack* may well be just what you are looking for.

Rating 8/10



The Megapack includes a vast selection of applications including games, utilities and office suites.

COBALT RAQ4

Nick Veitch takes a look inside everyone's favourite blue box

Price: £2,500-3600 **Web:** <http://www.cobalt.com>

Cobalt combined marketing savvy with the low cost customisable flexibility of the Linux operating system to create an astounding product with the original RaQ.

The concept was simple – create a device which, with the minimum of fuss, could be setup to serve multiple websites and web servers, with an easy to use web based administration service. Oh, and stick it in a nice eye-catching, small box that people would be proud to have in their racks.

In spite of heavy competition the RaQ series has proved very popular with service providers, mainly because of the factors mentioned above, but the RaQ is more than just a nice looking box that's easy to use:

it now includes some pretty cool tools, but more of those later.

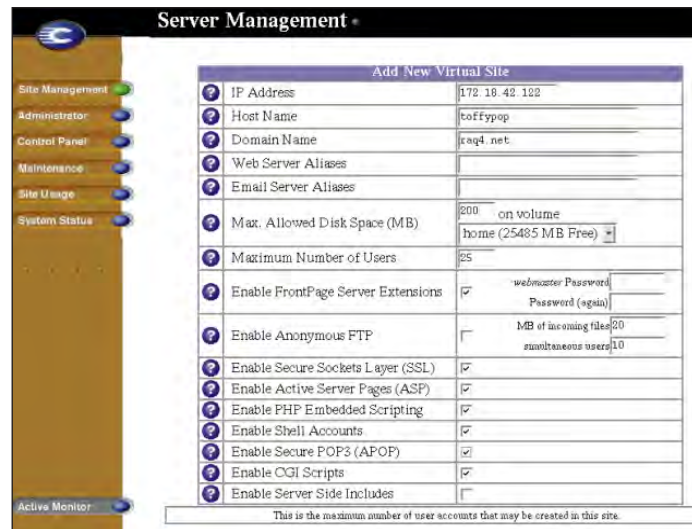
APPLIANCE

One of the fundamental concepts of the RaQ is that it's an appliance. It isn't meant to be anything other than a network box, and rather like a microwave oven or a fridge freezer, you don't want to be spending hours tweaking various settings to get it to cook your dinner, freeze the leftovers or indeed, serve your webpages.

Setting up a RaQ is therefore as easy as you could imagine. Physically installing it involves nothing more than connecting it up to the internal and external networks, via the clearly labelled RJ45 ethernet connections

at the back, plugging it in and entering some network information.

The front panel contains an LCD panel and some buttons with which you can enter your network



All of the RaQ4's management systems are accessed remotely via a web interface. It certainly make setting up a breeze.

specifications (IP address, gateway details, net mask etc). Then you can continue to configure the box completely via the network.

The device includes a set of scripts driven through an *Apache* web interface to change the configuration of the box. Adding websites is a simple matter of giving the site a name, setting up the disk space devoted to the site, and away you go.

Enable FTP to the box and you can upload your website from a local source. It's this speedy setup process that has made the RaQ series such a hit with ISPs – as Cobalt claim, you

HARDWARE SPECS

- 450 MHz Intel compatible processor (test machine included an AMD K6II/450)
- 2 SDRAM DIMM slots (PC-100) giving up to 512Mb RAM
- Dual Internal Ultra ATA hard drives supported
- Dual 10/100 Base-T Ethernet network interfaces Serial Console Interface
- External Ultrawide SCSI interface (40 Mb/sec)
- One PCI Slot for Expansion
- One USB port



Notice a theme? This is the first screen you'll see when booting up for the first time.



Cool, blue and very good looking. The Cobalt RaQ4 also works rather well.

could set up a host for a customer while they are still on the telephone.

Ongoing maintenance of the site is also fairly straightforward. Individual site administrators can change various aspects of their own site, including adding and removing users, constructing mailing lists and so on.

DEVELOPMENT

Part of the 'value-added' difference between the RaQ4 and earlier versions is the extent of pre-loaded software. Not content with the standard web plus cgi plus PHP server setup, the RaQ4 comes preconfigured with the *ChilliSoft ASP* software, allowing the development using Active Server Pages (for more



Various services can be made available from the Admin screen.

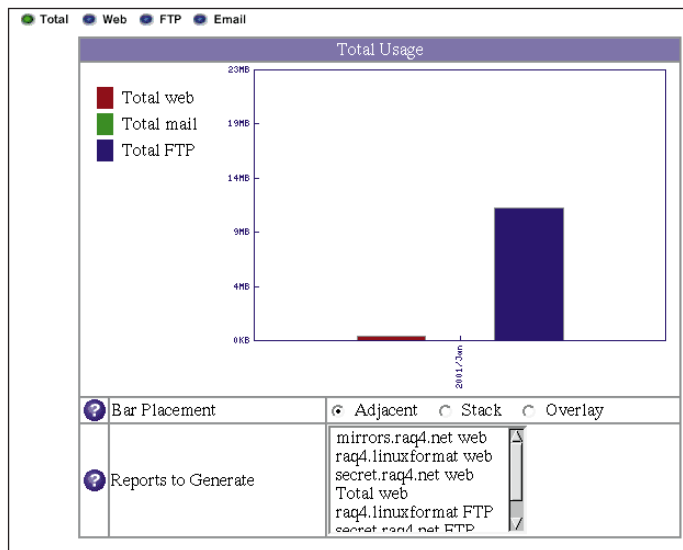
info on *ChilliSoft ASP*, see our separate review on page 34). There's also a full implementation of the Interbase SQL database and a 128-bit SSL layer preconfigured, to minimise the effort needed setting up a development environment.

As mentioned, Linux is at the heart of the RaQ server (and indeed, Cobalt's past employees include Larry McVoy and Dave Miller), so it is perfectly possible to add further extensions and developments to the environment, though these may not be manageable through the Cobalt admin server.

PERFORMANCE

The RaQ4 should be able to manage about 3.5 million web pages a day (though obviously this depends on the size and complexity of the pages), which is more than enough for most sites, even if you are hosting ten or so on each unit.

An early advantage of the RaQ was that it offered excellent



You can check up on web and ftp usage for the whole site. Bandwidth restrictions can also be easily imposed.

performance for a 1U box (when most comparable systems were 2U, or twice as high), which was a huge bonus for ISPs who could effectively double their performance/space ratio. The RaQ no longer has the 1U server market cornered: it now has to contend with machines from Dell, Compaq, IBM and even Sun (who recently acquired Cobalt). In some areas the RaQ 4 is not quite as close to the cutting edge as it once was. You don't absolutely need a super fast processor to run a web server, but you could serve many more pages or deal with more complexity with a faster Athlon processor (which don't cost any more now than the K6II did when the RaQ came out). Another bandwidth stumbling block is the drive interface. Although there is a SCSI interface, the internal drives are slower IDE models – cheap and easy to replace, but not the stuff of high performance dreams. Perhaps in this respect the RaQ doesn't offer as much value as some of the other server solutions, but currently the ease of use makes up for that in many users' eyes.

Web Usage by Browser (Top 2)			
Mozilla/4.74 (Macintosh; U; PPC)	54 (77.1%)	163k (50.1%)	
Lynx/2.8.4dev.8 libwww-FM/2.14 SSL-MM/1.4.1 OpenSSL/0.9.5a	16 (22.9%)	162k (49.9%)	

The package includes some useful if not extensive web statistic tools.

SECURITY

A cause for some concern might be the security aspects of using telnet. Sure, you don't have to switch it on, but it seems a little foolhardy when it could easily be replaced with a more robust remote shell tool like ssh. In fact, the only security feature mentioned in the manual is the security loop in the physical case of the machine.

But this device is aimed at ISPs providing http serving solutions – it isn't meant to be a firewall. While you could configure the RaQ to act as a gateway as well as a web host (and some of the tools required are already included, though not set up, on the hardware), that would be to miss the purpose for which it was originally intended.

MUCKING ABOUT

Although, to all intents and purposes, the RaQ is simply a webserving appliance, you can muck about with it to a great degree, without interfering with the admin configuration. In order to test its capabilities, one of our first needs was to get some web content.

SOFTWARE SPECS

- Linux 2.2 OS (2.2.14 in test machine)
- Apache web server, HTTP/1.1 compliant
- Virtual Hosting Services: name-based and IP-based
- CGI support
- Perl scripting
- PHP scripting
- InterBase 6 SQL Database
- ChilliSoft ASP
- SMTP, IMAP4, POP3, APOP email protocol support
- FTP, anonymous FTP access
- Telnet access
- DNS Server
- 128-bit SSL
- Front Page 2000 Server Extensions
- NTP client support
- Bandwidth Management
- Java Runtime Environment
- Security: PAM/Shadow Passwords
- Legato Networker client, Arkeia Backup Support

Connecting via telnet we compiled and installed *HTTrack* (see Hot Picks), and used it to mirror some popular Linux sites. Editing the crontab allowed us to update the mirrors nightly. Although you can do such things, a friendly message warns when you telnet in, that you could invalidate your warranty (make sure you back everything up!).

Although the RaQ may not represent such good value for many as it once did, the price tag is not excessive, and the prospective buyer will be getting an excellent web development unit, which is easy to use, flexible and reliable. To many professionals this is worth more than how fast the silicon goes. **LXF**

LINUX FORMAT Verdict

Ease of use: **10/10**
Documentation: **7/10**
Performance: **9/10**
Value: **8/10**

The RaQ continues to be the most complete web appliance on the market.

Rating 9/10

AGENDA VR3

John Waters was lucky enough to get his hands on this small object of desire...



From the iMac school of design. The Agenda VR-3 is nothing if not stylish.

Although it seems that Linux-based PDAs have been in the news for years, actually trying to get your hands on one is a little difficult. Take the Samsung Yopy for example: pictures have been circulating since early last year, but the actual device is still yet to see the light of day. The Compaq iPAQ is also mentioned in Linux circles, but it's little more than a curiosity at the moment.

It might come as a surprise then that there is a Linux-based PDA

already out in circulation. A US company called Agenda has built a unit called the VR3, and a Developer Edition was announced last year. A select band of developers – well those who were quickest to come up with the \$180 fee – were treated to the first batch a few weeks later. And of course, *Linux Format* has been in that select group from the outset.

FIRST IMPRESSION

The VR3 hardware is a very small,

very light, handheld device in a semi-translucent, almost iMac-like casing. The hard cover lid flips open like a Star Trek communicator to reveal the greyscale touch-screen. There are several buttons on the sides, and two on the front, with a cradle for serial port docking. An Infrared port promises compatibility with Palm devices and even consumer remote control options. Two AAA batteries provide the power.

That's the good news. The bad news is that, at the time of writing, the VR3 has a long way to go before it's consumer-ready. This is strictly a developer machine at this point: for example, activating the backlight requires a script launched from a command

line shell – hardly user friendly! The PIM applications are also at an early stage in the development cycle, so don't try and convince yourself this going to be a useful way to keep track of your personal life yet.

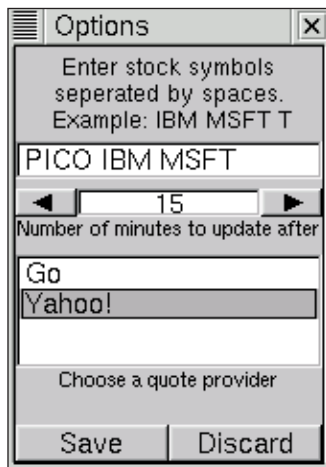
Our first impression with the VR3 was mixed. While the price is good, and having a real, live Linux handheld PC a joy, there were disquieting features. The initial release cannot be described as speedy. In fact, it's



The VR3 in all its diminutive glory.



ipaqlinux.com has some interesting stuff about liberating Compaq's handheld and giving it a Linux makeover.




The Agenda VR-3 can be used to check your shares...

something of a frustrating pig at times. Apparently this will be much improved as the software develops, but the processor seems to be struggling at times. The sound output is mono only which, in the MP3 age, is a shame. The wish list should include a USB interface, as serial syncs are so tedious with larger files.

Of course, the fact that not everything is perfect is the real beauty of this thing: if you sign up for the developer program you are getting in at the ground floor of a brand new, potentially very exciting, piece of hardware. The developer program isn't just for those keen to make applications: programmers are encouraged to make contributions to the core OS and applications too – with this type of open source policy, you can make a real difference to the final product.

A LOOK INSIDE


Looking for the technical stuff? Here we go. When you first power up the VR3, you'll see kernel boot messages scroll up the screen. Weeee! After a short time, the GUI will kick in, which of course is in the form of an XWindow server (XFree86). FLTK is used to provide all the buttons, menus and other controls used in the applications. As the CPU used is a 66MHz, 32-bit NEC VR4181 MIPS processor, you will need to set up a suitable cross-compiler on your system to join in the fun and games. The Linux kernel which the Agenda runs is a MIPS



...allow Minesweeper to overtake another area of your life...

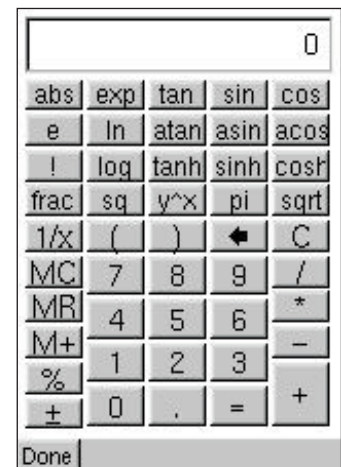
port that came out of the linux-vr project. In order to get newer kernels onto your device, you'll need to connect to your desktop PC by running PPP over the serial port. You can then use the monitor program on the Agenda device to copy across a new kernel and then flash it into place. Voila – a brand new OS on your device.

At the time of writing, the



...draw things like pretty flowers (ahh!) with the stylus...

developer mailing lists are incredibly busy, with a lot of excited programmers having a great time pushing the boundaries of hand-held computing. Help is only an email away: there really is a buzz around this little box. As long as you know the shortcomings, and are happy with this kind of system-level



...or budget for the developer version of Yopy.

programming, this could be the most fun you can put into your pocket. Being able to bring up a BASH shell while on the bus could just be the start of a new way to make sure you have Linux with you wherever you want to go... **LXF**

SMALL BUT PERFECTLY FORMED

Just like the Agenda VR-3, the Yopy Linux handheld has just had its developer release. The Yopy is fairly similar to the Agenda device, but – if the hype is anything to go by – should be something of a leap up the handheld evolutionary ladder.


The product aims to offer Internet access on the move, as well as advanced multimedia feature such as MP3 and MPEG movie playback, true colour video games, voice recording, digital camera add on and a geographical information system. Which is a lot for something that will fit in your pocket.

Visit www.gmate.co.kr for more information about Yopy and the developers kit which will set you back an eye-watering \$790.00.

Another interesting development is Compaq's iPaq PDA which, with some modification, is very Linux friendly. As part of their commitment to the open source ethos, Compaq are hosting www.handhelds.org which has a wealth of information about Linux handhelds from a variety of different manufacturers.

With most of the distro vendors dipping their toes into the embedded Linux market (see News), 2001 could well be the year Linux breaks out of the server and into your pocket.

The Yopy looks very exciting. Let's hope it's more than just vapourware.



WIN4LIN

Chris Howells discovers another piece of software aiming to help bridge the gap between Linux and Windows and dispel your dual booting miseries forever...

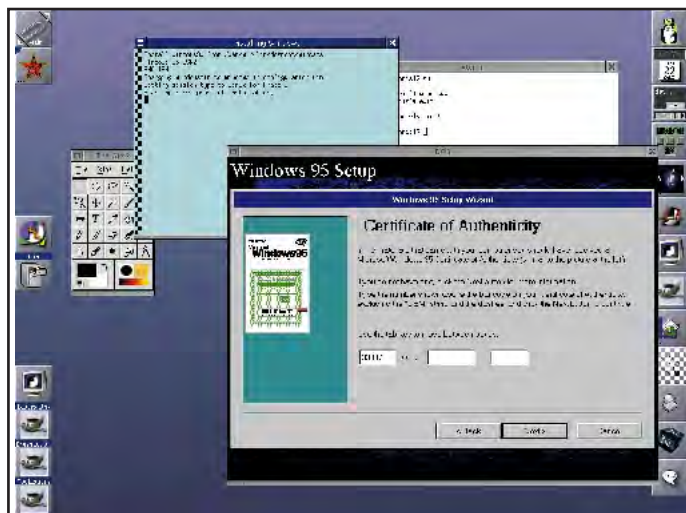
Publisher: Netraverse **Web:** <http://www.win4lin.com> **Price:** US \$89.99

Win4Lin aims to fulfil a goal which may be familiar to many readers – it allows you to boot an alien operating system – and use that OSs applications – under the GNU/Linux operating system, in a similar manner to *VMWare*, *Bochs*, or *Plex86/FreeMWare*. Unlike the similarly commercial *VMWare* however, which allows you to boot nearly any x86 operating system on an x86 box running Linux, *Win4Lin*, as the name suggests, is just limited to booting Microsoft Windows.

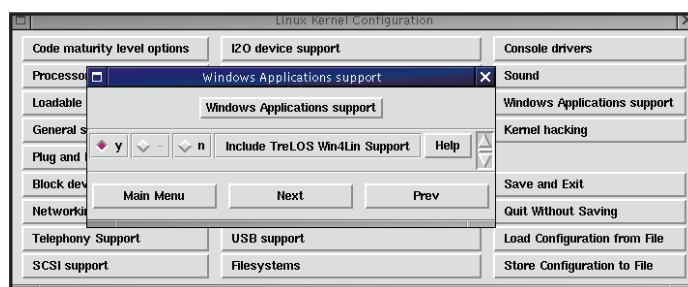
First impressions

Unfortunately though, in this second incarnation, *Win4Lin* is currently only capable of booting Windows 95/98. This means that if you are a corporate buyer looking for the capability to run Windows NT 4/2000 under Linux, then you will need to look at an alternative such as *VMWare*. Support for Windows ME is also likely to be patchy straight out of the box, due to the removal of the DOS layer in ME.

One of the dominant features when opening the box for the first time is the printed manual. This is of reasonable thickness, and is thankfully clearly written, covering all the vital topics such as installation, removal, and trouble-shooting. Another unmissable feature is printed on the envelope containing the CD-ROM. The strongly worded EULA (End User Licence Agreement) makes it clear that that *Win4Lin* is not



Setting up Windows in Linux is almost the same as setting it up in its own space. Just make sure you've got a valid licence.



Beware: Win4Lin doesn't like non standard kernels which means you may have to do a little manual configuration.

open source software released under the GPL or LGPL. This may upset some open source purists, but *VMWare* is similarly closed source, so *Win4Lin* is in good company. But that's enough of the politics – how does the software work, and more importantly, how well does it run native Windows applications?

Installation

Turning to the manual, there are no unusual system requirements, although you must have root (administrator) access to your machine. In addition to a *Win4Lin* licence you must also have your own copy of Windows 95/98 – a Windows licence is not included. If you received

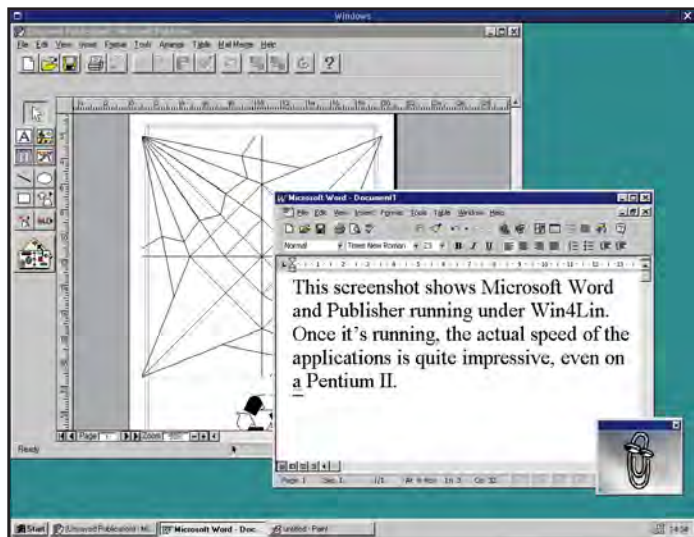
a copy of Windows pre-installed, then the OEM version should suffice.

Installation is comprised of two parts. First you must install the product itself, and then proceed with the installation of Windows from the Windows installation CD-ROM. Unfortunately though, the installation of *Win4Lin* is not one of its finer points. The first thing to note is that not all Linux distributions are supported. Debian anyone? Sorry, you'll be out of luck. Another thing that may upset some Linux users is that *Win4Lin* requires patches to the kernel to be applied, in order for it to function. Combined with the installation method, this is perhaps the biggest let down.

I proceeded to install onto a system running Mandrake 7.0 (kernel 2.2.14), but which I had upgraded to kernel 2.2.17 myself. If the distribution had still been running the stock 2.2.14 kernel, there probably would not have been a problem. It would simply have been a matter of installing an RPM package (containing the kernel image) with the correct distribution and kernel version. Instead, as soon as the *Win4Lin* installer detected I was running the 2.2.17 kernel, we received the alarming error message: "WARNING!!! You are NOT running the original Linux Vendor 2.2.XX kernel..."

For instances like this, *Win4Lin* supplies a kernel patch file which can be applied to the kernel source code tree with the **patch** command. The latest generic kernel patch on the CD was 2.2.16, so I headed off to the *Win4Lin* FTP site, and downloaded the latest 2.2.17 patch file. Despite some slight reservations, I successfully managed to apply this to the 2.2.18 source tree, and after a quick kernel re-compile, I was able to continue with the rest of the installation. Obviously, you will need to re-boot after installing the new kernel, so people trying to set an 'uptime' record may not be too keen on this.

Currently there is no support for systems running a 2.4.0test kernel, although to be fair, it would probably be too troublesome to develop the



Microsoft Word – complete with annoying ‘helpful’ paperclip – running in Windows, running in Linux

necessary patch while the 2.4.0 kernels are still rapidly developing.

You are then asked to copy the Windows installation files onto your hard disk. If your Windows installation CD-ROM is not bootable (95 isn't), you will also need the original Windows install floppy. After much experimenting with several Windows 95 setup disks, I eventually discovered one that the installer found agreeable, and continued.

The rest of the installation of *Win4Lin* and Windows is centred around the 'winsetup' program. Initially, it is only used to get the Windows installation under way, but

eventually it becomes the administrative centre point. You can then get onto the task of installing Windows itself. This is much like any other installation, although some sections of the installation are skipped, so during the installation you cannot configure the keyboard and language settings, since it defaults to US English. It is also not possible to customise the disk space requirements by choosing which components you wish to install.

After a short period of time, your virtual installation of Windows will reboot itself, and you will eventually be left staring at the Windows desktop. You are then free to do what you wish – or maybe not...

Running Windows

Initial impressions are favourable. After the slightly troublesome and long-winded installation, the speed that Windows runs at is very impressive. It feels nearly as fast as native Windows performance, and seems much faster than *VMWare*.

You can do nearly anything that you can under native Windows. As a test, I installed Microsoft *Office 97*, Norton *AntiVirus 2000*, and a few other things to see how they worked. To Netraverse's credit, the results were virtually flawless, and if you were not aware of *Win4Lin*'s presence, you would probably be hard pressed to tell any difference.

COMPILING A WIN4LIN KERNEL

This is not a comprehensive kernel re-compilation guide, but it should help if you discover the need to compile your own kernel to run *Win4Lin*.

First ensure that you have the kernel source tree, and have applied the relevant *Win4Lin* patch, using the `patch` command.

In an Xterm, switch to the kernel source tree, with `cd /usr/src/linux`, and then become root with `su`.

Enter the command `make mrproper xconfig` and then wait a moment. In the window which appears, configure the rest of the

kernel as you desire, and finally click on the 'Windows Applications Support' button. Ensure the 'y' option is checked, and then go back to the Main Menu. Click the 'Save and Exit' button. You should now be back at the command prompt. Enter the commands:

```
/make dep clean bzImage
modules modules_install
install/
```


Then configure your boot manager (LILO or GRUB) to support the new kernel, and you should be away!


However, it is important to note that since you are running Windows under a virtual PC created by software, a few things may be troublesome. And if you thought that *Win4Lin* would allow you to use your soft-modem under Linux, you're out of luck. Another thing that will not work is DirectX, the multimedia APIs which are crucial to virtually all modern Windows games. Disappointing. Sound does work however, as long as you have OSS drivers set up under Linux.

With *Win4Lin*'s configuration utility, you can define which serial ports and parallel ports will be available to your virtual Windows system. It is also possible to modify the bootup files such as `config.sys` and `autoexec.bat`.

Conclusion

Once *Win4Lin* is working, it works well. However, the installation is sufficiently awkward that some people may have problems working it, though anybody with enough knowledge to compile a new Linux kernel should not experience any problems. *VMWare* also requires a modification to the kernel to be made, however the installer is sophisticated enough to enable this to be done in an automated fashion, even if you are not running a stock kernel. If *VMWare* can manage it, there's really no reason that *Win4Lin* shouldn't as well.

However, *Win4Lin* currently does have two major advantages over *VMWare*. This is price and performance. The *Win4Lin* software is significantly cheaper than that of *VMWare*, which will please many home users. Also, as *Win4Lin* only focuses on running Windows, performance can be enhanced to a high degree. In the future, the capability to support Windows 2000 and ME would be nice, and this is likely to be essential with Microsoft's next generation Whistler operating system on the horizon. 



The Winsetup utility lets you get at most of the configuration options.

SUPPORTED DISTRIBUTIONS

The number of Linux distributions supported by *Win4Lin* is not exactly comprehensive. Those currently supported are:

Red Hat 6.x, 7.0
Caldera OpenLinux 2.2, 2.3;
eDesktop 2.4
SuSE 6.x
Mandrake 6.1, 7.0, 7.1

There also appears to be preliminary support for Corel Linux, so you may wish to keep a check on the Netraverse web site (www.win4lin.com) for announcements in the future.


LINUX FORMAT Verdict

Features: **6/10**
Ease of Use: **6/10**
Performance: **8/10**
Value for money: **7/10**

Once installed it works as well as can be expected. But work really needs to be done on the installation method.

Rating 7/10

MySQL



Publisher: New Riders **ISBN:** 0-73570-9211

Author: Paul Dubois **Price:** £38.99

Now that **MySQL** is available under the GPL, it has become far more accepted and widely used for ASPs and e-commerce sites. Fortunately, there are plenty of books on the market to help people get started.

Whether you're a hardened database administrator, or someone who's found out that their ISP supports *MySQL* and want to make an interesting web site, *MySQL* has something for everyone. Everything from compiling and installing *MySQL* (which quite strangely is towards the back of the book), to an excellent chapter on security aspects of running the system, something which is welcomed by many after the Apache.Org attack, is covered in depth, but with clearly written, easy to digest information.

As well as discussing the query structure of *MySQL*, there are also numerous chapters covering usage in a range of environments such as C and PHP, to Perl and shells. Quite conveniently, rather than just showing how to use a specific language to perform SQL queries, it covers the syntax and general usage of the languages, so if you've never touched Perl before, but want to be able to use *MySQL* data on a web site, everything you need to know is in this

book. There are no large lumps of code to bog you down, and each of the functions used is clearly described and commented.


In the preface it says: 'MySQL simply gives you the kind of information you won't find anywhere else'. While not entirely accurate, you certainly won't find this much useful stuff in any other book covering the subject. With the inclusion of a complete syntax, function and column type reference, there is little need to make the effort to go elsewhere for information. Fortunately, if you can't really understand the function syntax yet, there is a nice chapter covering the basics, such as how to rename a column or the varying ways to select data from tables.

If you don't yet have a *MySQL* book in your library, you'll be hard pushed to find an excuse not to get this one.

LINUX FORMAT Verdict

Once you have this book, you'll wonder how you managed without it, but you might want to wait for the next revision after 3.24 has been stabilised before shelling out.

Rating 9/10



PC HARDWARE IN A NUTSHELL

Publisher: O'Reilly **ISBN:** 1-56592-5998

Author: Robert Bruce Thompson and Barbara Fritchman-Thompson **Price:** £15.99

O'Reilly books have long been the staple diet for techies, covering a wide range of topics, not just Linux systems and applications but other operating systems (yes that one) and applications. However, they have not been known for producing hardware reference books, so this book starts to address this shortfall in their library.

For a book that is covering PC hardware it appears to be very light and perhaps an odd topic to cover under the series. When you compare this book with something like the classic *Upgrading and Repairing Your PC*, it is positively anorexic. However, this is a well written, thought out and logically formatted book which covers its topic in detail. There are chapters dedicated to each of the major components, such as memory, motherboards, hard disks and even power supplies. In addition there are comments throughout regarding upgrading older kit, and whether the upgrade is worthwhile.

Each of the chapters provides an overview of the component, types available, their differences and how to install the component into your PC. At the end of the book are two chapters dedicated to designing and building your own system. These chapters

provide a good reference should you be thinking of putting something together, providing example systems based on various budgets and requirements. The chapter on actually building a PC has a useful step by step guide – complete with pictures – on the installation of each part.


Unfortunately, when covering installations, the only examples that are given are for Windows based systems. This is not because the authors are negative towards GNU/Linux, it is mentioned favourably, however their expertise would appear to be limited to Windows based PCs. This does provide some amusing comments regarding Win9x and NT, such as: "...it would be worth spending the additional \$50... that NT costs", er, not convinced about that one. Nevertheless, this is an excellent book which should prove useful to all levels of user, providing a great deal of information.

LINUX FORMAT Verdict

An excellent book for both novice and advanced users, shame GNU/Linux is not covered in more detail.

Rating 8/10

The Linux Network Administrator's Guide



Publisher: O'Reilly **ISBN:** 1-56592-4002
Author: Olaf Kirch and Terry Dawson **Price:** £23.50

Almost all Linux systems are connected to a network in some way, either a LAN or a dial-up connection to the Internet. Therefore most operators of Linux boxes need some networking knowledge. The *Linux Network Administrators Guide* aims to supply that knowledge. This book has been around in one form or another for a long time, almost as long as Linux has been around in fact. I remember downloading a copy in 1993 and used it to set up my first Linux network.

This second edition of the book has been revised and updated by Terry Dawson. Considering the enormous changes brought about by the wider adoption of the Internet since 1993, I think that the book could have had a more radical re-write. Some areas such as Firewalling and IP Masquerading have been brought right up to date. There is a good history lesson and comparison of the three Linux firewalling systems ipfwadm, ifchains and, the most recent, netfilter.

The authors say that networking is such a wide subject that they couldn't include everything, but I do find their choice of subject matter a little strange. There is a chapter on


IPX and the Linux implementation of Novell's NCP filesystem, and also one on Taylor UUCP. However neither *Apache*, nor any other web server, gets a single mention. I'm sure that there are many more people aiming to set up a web server than want to use IPX, NCP and UUCP. There are also five chapters on Usenet News including details on setting up two different news servers (C News & NNTPD). Not many users want to set up their own News server. Also NFS and NIS are covered in some detail but there's nothing about Samba.

The material that is covered is clear and detailed with lots of examples, but before you buy it make sure that the topics that you are interested in are covered. If you don't want to fork out the £20 or so for the paper version the whole thing is available for download either in .PDF, LaTeX, Postscript and plain text versions from your favourite Linux Documentation Project mirror.

LINUX FORMAT Verdict

A detailed and authoritative reference, but it's beginning to show its age a little.

Rating 7/10



REBEL CODE

Publisher: Penguin **ISBN:** 0-71399-5203
Author: Glyn Moody **Price:** £12.99

This book, though it's title may not immediately suggest it, is an attempt to weave an entertaining and informative story around the history of Linux and the Open Source phenomena. Although the book does take in other aspects of the open source community, the main thread centres on Linus Torvalds and the developing story of the operating system he created ten years ago.

The book is certainly well researched (the author claims it's the result of over fifty separate interviews), and in fact, in some places it seems to be a little too detailed. There is a tendency in the book to back up information by including the exact email messages sent to usenet groups, then analysing the intentions of their author line by line. This is fine for quite a few important moments, but is a little overdone in parts.

Apart from this the book is generally well written and an effort has certainly been made to create an enjoyable read as well as a comprehensive history of the GNU/Linux project. You'll discover plenty of things you didn't know, and a good effort has been made to analyse the methodology of projects as diverse as the Linux kernel, the Mozilla project and the Debian

distribution and explain why certain things worked out the way they did.

While the book often returns to Linux, there are plenty of forays into the rest of the open source community, including a look at the organisation which in many ways allowed Linux to exist, the Free Software Foundation.

The Linux world moves at such a fast pace that no print-based book can hope to cover the very latest developments, though *Rebel Code* surprisingly up-to-date. Despite offering a few pointers for the future of Linux, this is, essentially, a tale without an ending, and though you may quibble occasionally with the interpretations the author implies in a few places, it is about as comprehensive a book as you could wish for. It may not be an 'unputdownable' page-turner, but it is very readable, and highly recommended. **LXF**

LINUX FORMAT Verdict

An enjoyable read for anyone interested in the history of open source development. Though the author can get a little bogged down in his quest for 'completeness'.

Rating 9/10

THE BEST NEW OPEN SOURCE SOFTWARE ON THE PLANET

HOT PICKS

This month's Hot Picks has gone a little Internet mad, with some font hysteria thrown in.

If you're looking for the best new and improved open source software the planet we call Earth has to offer, then you're in the right place. Every month we scan the new releases, and bring you the best of open source, whether that means exciting new projects, or radically updated apps that

you simply must have. This month we noticed quite a few great internet-related tools that we think worthy of the accolade "Hot Pick". You'll find all the software mentioned here on our CD this issue, so enjoy! And if you have any suggestions for future issues, mail us!

QUANTA PLUS

Version: 2.0b4 **Web:** <http://quanta.sourceforge.net/>
Source: RPM/Source

If you've ever wanted to produce your own website, you'll have run into a bewildering range of tools and technologies. While it's true that you

can write simple HTML using a text editor like *vi*, you'll probably want to add one or two images (which will need an editing tool like *The GIMP* or

Photoshop to prepare). Then you might want to add style sheets, possibly some JavaScript widgets to make those nice bevelled buttons change colour when a visitor mouses over them, and so on. Then you discover that you're handling several

dozen HTML files, and a simple text editor won't do the job any more; you need some way of visualising the structure of your website.

Enter *Quanta Plus*. This is one of a range of Linux-based HTML editors that are sprouting website management tools like the ones that made Microsoft *Frontpage* so popular (despite the execrable quality of the HTML it produced). *Quanta Plus* release 1.04 for KDE 1/1.1 is currently stable; *Quanta Plus* 2.0, targeted at KDE 2, is in late beta and thoroughly usable already. To install *Quanta* 2.0b4,

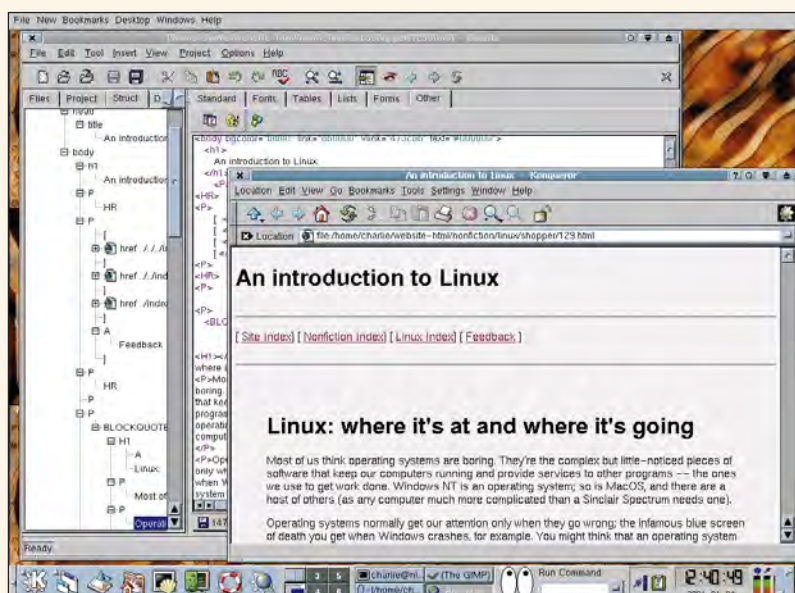
you'll need to go through the chore of compiling it yourself: download the tar.bz2 file from quanta.sourceforge.net, decompress it (**bzip2 -d <filename>**), unpack it using **tar**, then type **configure && make** on the command line inside the *Quanta* directory.

START A PROJECT

When you first run *Quanta*, you'll see a multidoc window with numerous tabs. The big window at the right is for editing or previewing files; the one at the left has options for viewing the list of open files, a project's structure, the tag structure of the current HTML document, and the online documentation. (It follows the general style of the *KDevelop* integrated development environment by providing a library of reference docs that you can browse.)

I took my personal website (roughly 300 HTML files and 40 images) and created a *Quanta* project from it, using the Project menu. A project is a bunch of files that *Quanta* can work with as a collection and upload to a web server; you can import and work on files from the web, from a local directory, or from a CVS repository.

The editor is pretty much what you'd expect: it provides a colorized view of your HTML source, along with various tools for inserting tags. It's not



You can preview your creation either within Quanta or through your browser, but you have to go back into the code to do any editing.

a WYSIWYG system, and although you can preview the appearance of your HTML you can't edit the preview. One interesting feature is the 'structure' view in the left-hand pane, which provides easy navigation and selection in the document. Another is the ability to edit tag attributes, including scriptable events, and to edit and attach style sheets for your documents. The editor is also customizable, and you can redefine the key bindings to suit your own preferences if you feel like it.

ON THE INSIDE

Internally, *Quanta* is an XML editor; it uses an XML DTD that specifies the attributes and contents of tags in a document, and which in turn drives the syntax colouriser, structure view, and so on. This is a promising approach as it permits *Quanta* to rapidly be updated to support different levels of HTML and different highlighting modes – as an editor it can provide syntax colourisation for HTML, Perl, Sather, Python, PHP3, JavaScript, LaTeX, and a whole range of other file types already. Hidden in the functionality, you can add your own programs to the configurable toolbars: the Options->Configure Toolbars dialog is flexible enough that you can add an external script or program, and send it the current selected text (or entire document) as its input; the output replaces the appropriate chunk of text. You can even use this to define chunks of boilerplate text, which can be inserted with the click of a mouse button on a toolbar. Thus, *Quanta* is easily extensible to do tasks that the designers may never have envisaged.

Quanta Plus isn't perfect. There are a number of desirable features missing from it; perhaps most disappointing of all, the project view doesn't allow you to edit the structure of a project by dragging files between directories and automatically updating their links. On the other hand, the software isn't finished yet, and those features it provides are polished and professionally implemented. This is well on the way to being the best web editor for KDE.

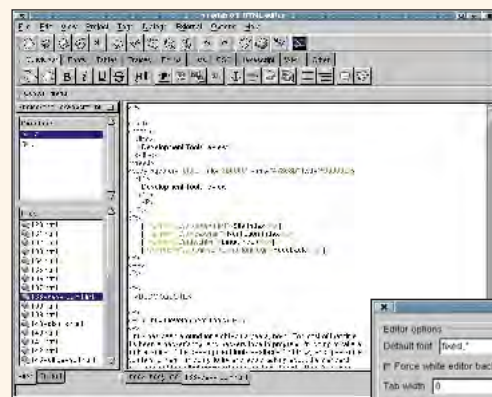
BLUEFISH

Version: 0.6 **Web:** <http://bluefish.openoffice.nl/>
Source: RPM/Source

Bluefish is a different HTML editor for Linux. Unlike *Quanta*, it isn't tied to a specific desktop environment – all it needs is X11 and the GTK widget set. If you have an up-to-date *GNOME* desktop environment, all the libraries you need to run *Bluefish* are present – but it isn't a *GNOME* application as such, and doesn't need all the *GNOME* extras.

easy to setup a table for text layout and adjust all its parameters (such as spacing, alignment, background colour, and so on). There are some predefined JavaScript widgets that can be inserted with a mouseclick, and virtually all the HTML tags you could want are accessible from menus or toolbars, and come with a dialogue that offers you all their acceptable options.

On the minus side: project support in *Bluefish* is incomplete. There's a project menu, but some essential items (such as 'create new project') are missing.



Colour coding improves the legibility of your code.

Like *Quanta*, *Bluefish* presents a multidoc view of your work; edit pane at right, directory window at top left, and list of files at bottom left. However, *Bluefish* makes no pretence at WYSIWYG, even to the extent of providing a non-editable preview mode; to preview a page you need to feed it to an external browser (such as Netscape) via the appropriate menu option. One minor annoyance: neither application seems to support preview options for different browsers – vital when you're developing content that needs to look good on *Lynx* as well as *Mozilla*. Although you can probably work around this.

EASY WRITER

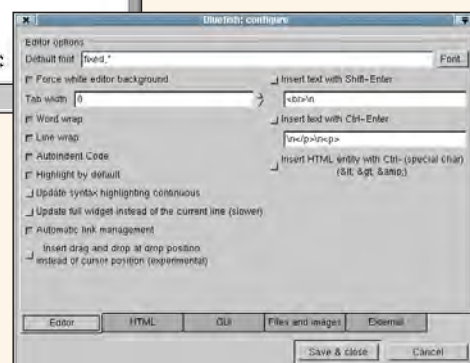
All the basic HTML tags are available for insertion. There are some nice touches: the table wizard makes it

and an External menu that you can hang commands and filters off – commands are executed and their output inserted into the HTML, while filters run on the selected text. (It's preconfigured to hook into *Weblint*, a good idea when dealing with an editor that doesn't understand the structure of an HTML document.) *Bluefish* also has a set of menus and tools for editing WML, the language used by WAP phones: this is the only WML editor I've seen on Linux so far. It also has a CSS editor that looks rather more comprehensive than *Quanta's* – however it suffers from a lack of contextual help: if you don't know what a keyword or parameter means, you're out of luck, and you'll need to go rummaging through the CSS or HTML standards documentation (wherever you left it).

YOUR ON YOUR OWN!

It's in the arena of documentation that *Quanta* shows an edge over *Bluefish*. *Bluefish* is simply not documented; although *Quanta's* documentation and help is incomplete, at least there's a

documentation browser to provide a framework for adding it. When you're writing files using a structured format like HTML, CSS1 or PHP, you really need a language reference to hand. *Bluefish* is still a work in progress, and it seems that documentation is way



You can configure almost every aspect of the program.

And it's not a structure editor; double-click on a table tag to select it, call up the "table" dialog, and what you end up with is a new table tag (with whatever parameters you specified) inserted within the text of the previous one. You have to understand HTML in order to get a valid web page out from here; this is not an idiot-proof editor.

On the plus side: it seems to be easier to add external programs to *Bluefish*. There's a Custom menu that you can do whatever you like with,

down the priority list. This is a common problem of open source projects, and it is to be devoutly hoped that developers learn to pay more attention to it.

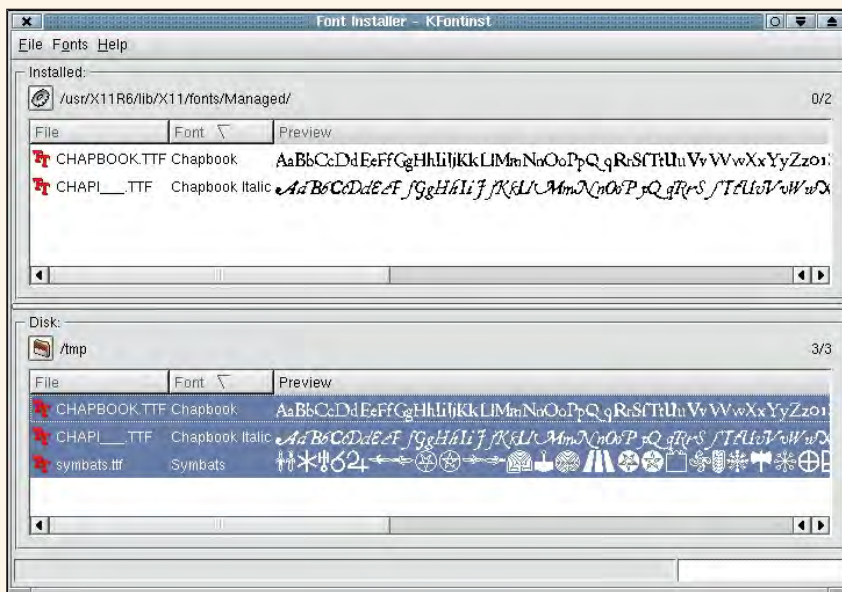
So why, then, do we bother reviewing *Bluefish*? The simple answer is: because it shows a lot of promise. Add documentation and point help, fix the rough edges with project handling, and it'll be a powerful and productive HTML editor. Being GTK based, it can be turned into a *GNOME* application relatively easily. And with a little bit of work it will rival *Quanta* as an HTML editor.

KFontInstall

Version: 0.9 **Web:** <http://www.cpdrummond.uklinux/kfontinst/> **Source:** Source

One of the most annoying problems with Linux is nothing to do with Linux itself – it's the difficulty of installing and using fonts on the X11 windowing system. X11's font handling is like something out of the 1980's. You can plonk a bitmap font (or – shock horror! – an Adobe Type 1 font) in a maze of twisty little directories, run various utilities to update index files so that X11 knows it's there, re-start your X session, and see the font on screen. Except that when you try to print it won't be there, because WYSIWYG printing on Linux goes via the *Ghostscript* postscript engine; you also need a copy of the font installed where *Ghostscript* knows about it, and you need to tell *Ghostscript* that it's installed. And then when you run *AbiWord* or *StarOffice*, they don't know about the new font, either: because you haven't told them you installed a new font.

The picture gets a little better on newer systems. *FreeType*, a freeware TrueType font interpreter, provides a



KFontInstall lets you see which fonts are available, and which are already on your system. Very useful.

thing called a 'font server'; this takes TrueType fonts and transforms and scales them into a bitmap that X11 can understand. (The latest release of XFree86, 4.0.2, goes further; adding facilities which should ultimately result in an X11 display server with built-in TrueType font handling and support for anti-aliased fonts – the sort of thing MacOS and Windows have had for a couple of years now.)

KFontInstall is an attempt to fix some of the hideous problems of administering fonts on a Linux system. Given a Redhat 7 or equivalent Linux

installation, *KFontInstall* will (with a bit of configuring) allow you to install and uninstall Adobe Type 1 and TrueType fonts easily. The fonts will be

This is where you set the paths for general and application specific fonts.

available to X11, KDE, *Ghostscript*, *StarOffice*, and (with some manual hacking) *AbiWord*. *StarOffice* support is of most interest here as it is famous for its 'difficult' font support.

The most important thing to understand about *KFontInstall* is that although it provides a nice user interface with two panes (one showing a preview of available fonts, and the other showing the current installed and managed fonts), it takes some setting up.

READ THE FINE MANUAL


Before running it for the first time you really need to read the documentation. There are directories to create, one or two configuration files to edit, and so on. In particular, if you want *KFontInstall* to correctly install fonts for *StarOffice* you may need to identify the printer driver description file that *StarOffice* is using and manually add a line to it. This is not a point-and-shoot tool for taking

all the pain out of managing fonts on Linux; however, it enables you to experience the pain just once, instead of forcing you to jump through the same flaming hoops every time you install a new font.

KFontInstall makes the X11 font side of things easy: you create a spare directory for the managed fonts, tell the font server about it by editing a configuration file, and thereafter *KFontInstall* gives you the ability to add fonts to X11 in a point-and-shoot manner. *Ghostscript* is harder, mainly because some versions of *Ghostscript* understand TrueType fonts, but others don't, and the application also wants its fonts installed somewhere else, requires font metric files and insists on a pile of different configuration files. Still, if you follow the instructions, *KFontInstall* makes automatic *Ghostscript* configuration reasonably easy. The major headache is *StarOffice*, which has its own font rendering system. You will need to identify a number of *StarOffice* configuration text files and edit them by hand before *KFontInstall* can install fonts that *StarOffice* will happily use. It's fortunate, then, that the manual for *KFontInstall* is comprehensive and gives a reasonable explanation of what's going on here. But it's still not suitable for amateurs – who could easily render *StarOffice* inoperable with a few rash keystrokes.

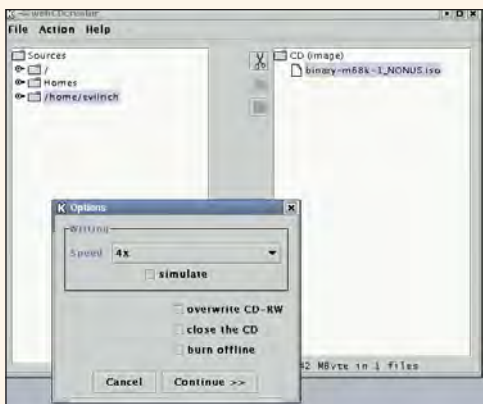
RESULT!

In summary, *KFontInstall* is a valiant attempt at tackling a very hard problem – that Linux doesn't have a standard way of installing and managing typefaces. It's by no means perfect, but it's better than doing everything by hand. Hopefully some of the Linux distributors (SuSE and Redhat, are you listening?) will take note and incorporate a pre-configured copy of *KFontInstall* into their future releases. That would be very good indeed from the point of view of a user who just wants to drag a TrueType font file onto an application icon and have everything work – which is how it should be.



webCDwriter

Version: 1.0rc3 **Web:** <http://www.uni-bielefeld.de/~jhaeger/webCDwriter/>
Source: RPM/Source



WebCDwriter can be accessed from any machine running a java enabled browser.

Picture the scene. Everybody in your office needs to be able to burn CDs now and again, but not so often that it's worth investing in a CD burner for each machine on your network. The solution would be to somehow share a single CD burner amongst your users, but how? Well, forget about mucking about with NFS and remote logins, *webCDwriter* is a package which makes the whole thing easy. And client machines don't even have to be running Linux. Any machine running *Internet Explorer* or

Netscape or with a Java runtime environment installed will do.

The core of *webCDwriter* is the CDWserver, a daemon which runs on the network host with the CD writer drive attached. This is basically a network frontend to the standard CD writing tools, *mkisofs* and

cdrecord. Client machines communicate with this server using a special protocol and tell it to create and write CD images.

The user accesses the CDwriter through the other half of the *webCDwriter* package, a Java program called *webCDcreator*. This is a GUI-based application with which the user remotely instructs the CDWserver to build ISO images, create audio tracks, and burn CDs. *webCDwriter* exists as both as a Java applet and a Java application. Using it as an applet

means you don't need to actually install it on any client machines. You can let users access it from a web page displayed on their browser and served up by a copy of *Apache* installed on your CD server. The GUI requires *Swing*, so you will need Java 2.0 and the Java browser plug-in installed on the clients.

The user starts a new CD session by selecting the type of CD he or she wishes to create. This can be a data CD, an audio CD or you can write a locally-stored ISO image to disc. Unfortunately there is no support for mixed-media discs, but that may be something for the future. To build a data CD, you simply select the files and directories from your local filesystem that you wish to be stored in the ISO image. You then hit burn, give the disc a volume name, select whether you want to create the image with Joliet or Rockridge extensions or as an Apple HFS image. Hit continue and all the files will be transmitted to the server where an image will be built and stored. (You can also configure the CDWserver to

build ISO images on the fly as it is burning the CD.)

When the contents of your CD-to-be have been transferred to the server and, if necessary, an ISO image prepared, the system will try to obtain access to the CD writer itself. Many users can create ISO images at the same time – providing you have enough disk space on the server – but obviously only one CD can be burnt in a drive at a time. The server will keep the drive assigned to you for a defined period (the default is ten minutes) in which time you'll have to walk over to the machine and put a blank CD-R in the drive. If you don't, your session will time out and the next user in the queue will get the chance to burn their CD.

webCDwriter has all the usual controls you would expect of a local CD writing package. You can select the write speed, there's a test mode, and you can choose whether to fix the CD after writing a session. *webCDwriter* also supports CD-RW drives with an option to overwrite CD-RW discs.

This is a well thought out and well implemented package. Okay, the creator GUI for building your ISOs image is a bit spartan – with no drag and drop – but it gets the job done. We've been using *webCDwriter* in the *Linux Format* offices for about a month now, and we've found it both reliable and stable and very easy to use. It's a great way to avoid duplicating resources unnecessarily.

Downloader for X

Version: 1.20 **Web:** <http://www.krasu.ru/soft/chuchelo/> **Source:** RPM/Deb/Source

As Linux pretty much originated on – and because of – the Internet, it isn't surprising that the Internet is often the first port of call for Linux users when they want to get updates, patches, documentation or just about anything else.

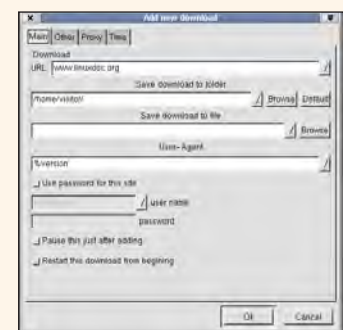
The trouble is, you don't want to be hanging around while downloading

the latest kernel source, you want to be doing other things, without the need to watch over the download process. And if you find a site with the right documentation on it, you might want to download at least part of the site, rather than just save an HTML page from it.

For these occasions – and many

others – *Downloader for X* becomes an invaluable tool.

We have covered this software before in these pages, but the latest version is such an improvement, that it has become the first project to be featured in Hot Picks twice! The main improvements are a number of bug fixes, but *Downloader for X* is now far



Select the source and destination and hit 'OK'

more stable, and handles a greater range of servers and proxies without any problems. Also the drag and drop feature has been tweaked, and ➔

REVIEWS | hot picks

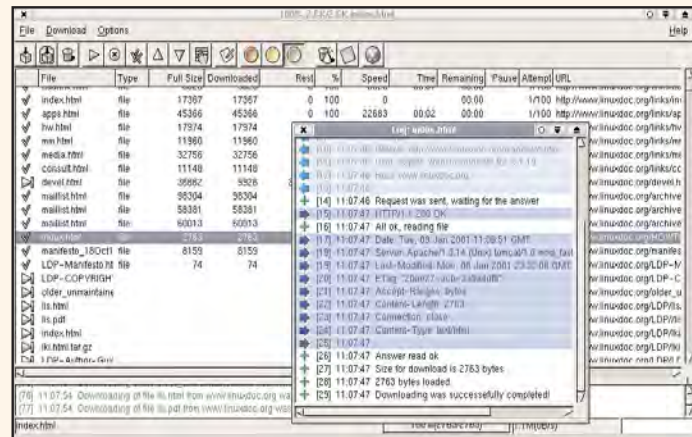
➔ now works correctly with Netscape 4.7 - I guess getting working with Netscape 6/Mozilla will be next on the list.

There are also huge improvements in the tools for downloading complete ftp sites, which avoid a lot of the duplication encountered previously, and the HTML parsing is more thorough. Also, more language support has been added, so the software is now available in 16 languages!

As with the previous version, it is possible to define custom limits with regards to recursion or sticking to

particular servers. You can create custom host lists with defined characteristics for each host, supply usernames and passwords for protected domains, and limit recursive searches to particular types of file - particularly handy for retrieving things like sources from an FTP site, or particular rpm builds.

This is one of those rare programs that stands head and shoulders above the competition in the flexibility and reliability stakes. What it does, it does to a very high standard. For complete mirroring of sites, see our next pick!



Downloader for X keeps an extensive log of each file transfer negotiation, so you can see where things are going wrong (or right).

HTTrack

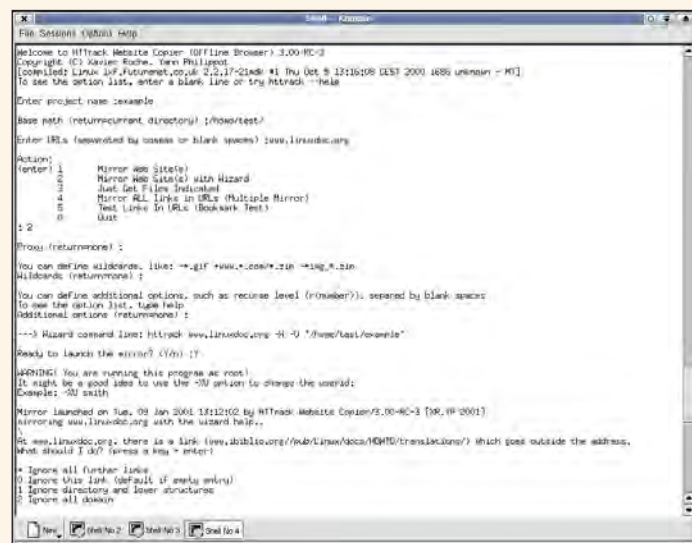
Version: 3.00 **Web:** www.httrack.com **Source:** RPM/Deb

This console program is designed for one purpose only: mirroring websites. If you run a mirror of a website, or if you want to make your own personal mirror of a webpage, websites or ftp directory, this is one very useful bit of code.

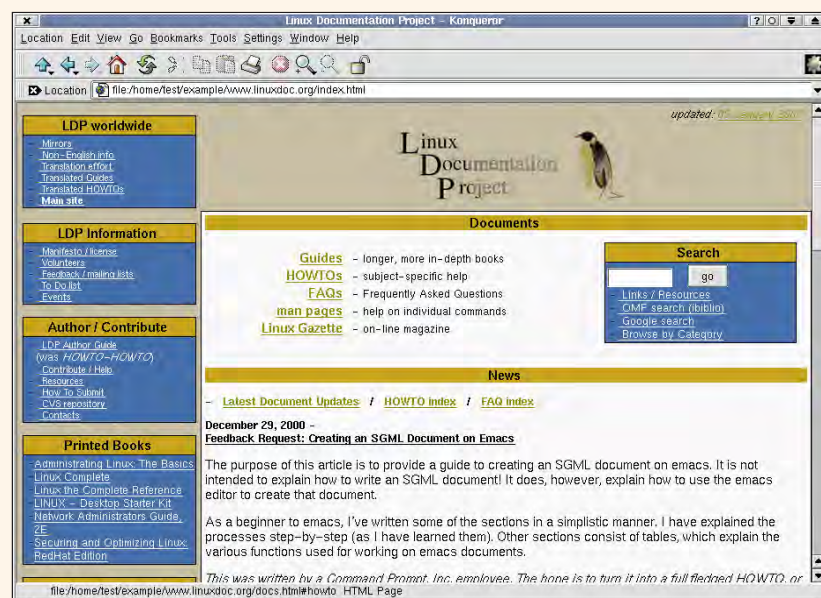
Although this is strictly console driven, there are a variety of ways to use it. If you run the program with no arguments at all, you'll get an

interactive mode, which asks questions such as what site you wish to mirror, where it should be stored and so on.

Obviously, because it can be run from the command line, it is an easy matter to set up a cron script to automatically update the mirrors over whatever period you choose. You can still supply additional parameters such as recursion depth



As is common, the console interface hides an extremely powerful and useable application.



We used HTTrack to mirror a number of sites, including The Linux Documentation Project, for use in the RaQ4 review.

and proxy preferences, and there are a variety of structure methods to use (which change the way the local mirror files are stored). If you choose not to download the entire site, you'll find that the links on your local HTML version will still point to the remote original, so the structure is still intact, even if it isn't all stored locally.

Having said that, it's still easy to use. For example, if you wanted to mirror www.linuxdoc.org. All you have to do is **cd** to the directory you want to

store the mirror and type: **httrack www.linuxdoc.org** and away it goes. This uses the default mirroring options which will not stray from that domain, and will fetch all links.

The software will generate a front page for the directory where the mirror is stored, with links to all the individual sites mirrored. A cache is kept to allow the sites to be updated more efficiently, and to resume incomplete transfers. There is also a DNS cache, to make further retrievals that bit faster.

As well as using this software to mirror existing sites, it can also be used to effectively check sites for broken links. Comprehensive logfiles are generated which can be used to find broken links and other potential problems with the site. **LXF**



Welcome to the *Mailserver*, the place where the most important people in the Linux community get to write in and tell us what they think about anything Linux related. Who are these people? Why, they're you, of course.

We enjoy receiving all your emails and letters relating to the magazine and all matters Linux. We don't mind you having a rant about things - we actually enjoy it (though we may not agree with you). More than that, it gives us a very valuable insight into what sort of magazine you want to read. Nobody has ever accused a Linux user of being shy or short of an opinion on any given subject, so why not share your views with us, and indeed, all the other thousands of Linux enthusiasts who read this mag.

The only thing I would ask you is that you include your name, and a rough geographical location at least. It may be obvious to you who you are, but we like to know who our mail, particularly the email, is coming from!
Nick Veitch
Editor

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

Putting the "hard" in hardware...

Having just purchased *LXF10* I was interested to read J. Faulk's letter with respect to newbie stuff. Being one too, I have also found the 'obvious' to be anything but on some occasions!

I've been using various M\$ 'operating systems' (I use the phrase loosely) for many years so I know what to look for when they (often) go wrong, but with Linux it's much more fun! I still don't fully understand the directory structure since I've only been using SuSE 7.0 for a short time. Hence, finding an offending file can be a tricky business, not to mention deciphering some of the alleged fixes on the net.

Which brings me neatly to the point of my letter - I notice you have a 'practical guide' to getting 3D hardware to run properly planned for the next edition. I am in the (unfortunate for a newbie) position of owning a Nvidia TNT2 Riva M64 graphics card, which is not the most straightforward device to get to work fully under Xfree86 v4. Could we possibly have some very basic instruction in how to get NVidia based cards to

work? I have been assured they will work and have tried a variety of fixes with no success (to date) from the SuSE and NVidia web sites, but then I'm just a newbie...

LXF is the best mag I've found for introducing Linux to the masses and sorting out problems. As a mark of this achievement I'm going to subscribe! Nothing particularly amazing in that you may say, but the last computer mag I subscribed to was *Dragon User* way back in the eighties! Please keep up the good work and last longer than *Dragon User* did!

Kind regards
Matt Oakley, Leicester

Well, hopefully you will find all the info you need to get your card going in our feature this month. It isn't just newbies who have problems with 3D hardware I can tell you. Part of the problem is that often the manufacturers details on various bits of hardware aren't specific enough to help you set up the X server properly. And then when they release drivers themselves, often these aren't sufficiently documented.

If there are any other areas of Linux that are giving you problems, feel free to write in to our Linux answers section or why not suggest a tutorial or feature!




UKLinux.net provide solid Internet access provision without the annoying Windows bias.

Penguin friendly ISPs

I'm looking for a cheap ISP - either one with a monthly subscription but no call charge, or one with local call charge but no monthly subscription. The one I'm using at the moment was worth paying for three years ago, but it's service standard has deteriorated recently and I don't think it's worth paying the premium price any more.

Do you have any advice about which ISPs are Linux friendly? The basic problem is that many seem unable to register users except by running a CD-ROM which requires to be run under Windoze. Their, so called, technical support does not extend to being able to tell you what phone number and PPP authorization information you need to configure, nor any other technical data which may be useful e.g. the IP address of their domain name service or the name(s) of their POP and SMTP servers.

As a last resort I suppose I could install Windoze on an old box, then install from a CD, dig out the required config info, set up on Linux and throw Windoze away again, but it all seems a bit of a bind! Surely there must be services somewhere which allows you



Configuring 3D hardware is something of a dark art with Linux. Fortunately, we have out best dark artists working for you.

register on a web page and provides the necessary configuration data in a raw form?

Regards
Brian Beesley
via email

Two Linux friendly ISPs we know of are www.uklinux.net and Scotfree-uk.net, try them first....

Disgusted of the net

I was just browsing my spam box, catching up on some news, and I noticed the header on the latest *Linux Format* newsletter.

I was *disgusted* to find that the X-Mailer was a Microsoft product. If you are making a Linux magazine, or news letters, you should use Linux software. It is not like Linux does not have good email clients (read your reviews issue 3, July 2000). I hope you are using Linux to produce the magazine with?

Back when Future Publishing was making *Amstrad Action*, they used to claim that they produced that using the Amstrad CPC range. Don't even attempt to claim that it's not possible to create the magazine using Linux software. You could use many office products (Roundup Issue 9, Dec 2000), LaTeX (What on Earth issue 1, May 2000), or even use html, and convert it to postscript.

One other thing: you produce a mag for the open source community, please would you supply back issues

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

in pdf/ps/gz format. I guess you don't want to send out the current issue so people need to buy it if they want to read. But if there are no copies left to order, please can we have the chance to read it anyway?

Adrian Moye
via email

Yes, sorry about that. But, we do live in a real world. This company produces a great number of magazines, and a lot of the systems here are similar across magazines. Mr Fenwick, from whom the message originates, is involved in many Future websites, and I don't expect he runs a different mail client for each one.

We do take Linux seriously, which is one reason why our website runs on a Linux server, whereas all the other Future websites currently use Windows based servers. But there are limits to how far this can reasonably go. In terms of the production of this magazine, pages are laid out on Mac's running Quark Xpress. There are a few very good reasons for this.

1. Every other magazine produced here uses the same system. I don't think the IT department or the company accountant would relish the idea of installing separate production systems for every magazine.

2. I would actually contest the point that professional level publishing solutions exist for Linux – TeX is a really great piece of software, but I wouldn't want to try and lay out this magazine using it. Assuming that they did, who would use it? One of our professional designers who has spent all their life using a Mac?

3. Would you expect the Official Playstation magazine to be laid out on Playstations. On the editorial teams, we run Linux all the time. We don't even dual boot. We even offer help to our IT department when it comes to looking at Linux products. But we are never going to get, for example, the scanning department, to scan in pictures on Linux machines. Why would they want to when they have a system that works perfectly well (for them) already?

I can see what you are getting at, but Linux isn't just about disliking other systems. I think one of our machines here has a Microsoft mouse – should we chuck that out the window too? If it makes you feel any better Mac Format, Edge, PC Answers, PCPlus, the official Playstation magazine and a whole host of other magazines use Linux servers for file sharing and printing. Why? Because it's a damn site better than the server they used to run on...

In answer to your second question, we are hoping to make some material from issues which have sold out available as PDFs on our website and CD in the future. Watch this space.

Tough going

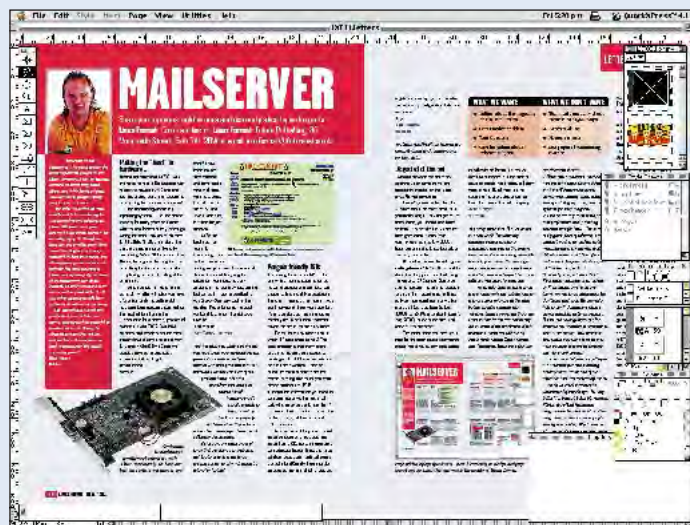
About 12 months ago I had my first encounter with Linux, it was SuSE 6 with YaST 1. I never finished an install with it because I was not confident with all the necessary partitioning. Then came Corel Linux and I tried more than once to install it in a Windows folder, but it never ran no matter how I booted to it.

About this time I decided to reconfigure my PC which was running a 1.7Gb disk as C: and an 8.4Gb as D: so I reinstalled everything on the 8.4 (now C:) which gave me the 1.7 to play with in Linux.

I went for Red Hat 6.2 which is actually more complicated to install. I also chose the GNOME desktop option, and it took four attempts at installing before I saw the little check box to allow graphical logon. Who cares what you can do in a console or at a prompt!

Once up and running, Internet configuration was easy and so was printer set up, but it crashed more than Windows ever had, mainly this was caused by Netscape, but not exclusively. So, I tried the KDE install and it was marginally better. Then along came SuSE 7. YaST2, it said in your magazine, was so easy, so let's go and replace Corel. Oh poo! Yast2 and my PC, possibly my video card do not get on so it's the dreaded YaST1 but we get there and I like SuSE 7, and it can use my sound card, at least some times and has *Star Office 5.2* and some other nice bits of software.

But wait, there's more, along comes Corel Linux 2nd edition. No KDE2 but it has the Corel File-manager, which allows access to Windows C: drive, my favourite ability, and a graphic lilo with all sorts of options, but no sound support. ➔



Page within a page spookiness. There is currently no design and page layout app for Linux that can match the quality of Quark Xpress.

LINUX FORMAT

To have your say

please contact:
LINUX Format
Future Publishing
30 Monmouth Street
Bath BA1 2BW
email:
linuxformat@futurenet.co.uk

LETTERS | your views

➔ It now finds printers easily, has the easiest Internet wizard I've seen and imports all *Internet Explorer 5* settings and email settings into *Communicator*. Its Archive installer leaves the others for dead, helpful if you're a Newbie, as only *Star Office* comes ready to 'Install'.

All this is well and good but no sound support, a scanner still out in the cold – despite SANE seeming to be installed and running – and a digital camera that's Windows only for the foreseeable future. It hasn't cost me a lot, but it's not killing off Windows just yet.

I'm typing this on *Abiword* (second attempt because that just vanished and only the restart button would work), I have *Corel Wordperfect 8 for Linux*, *Star Office 5.2* and there is *Kwrite*, so word processing is covered, the rest of *Star Office* seems OK, *Compupic* is handy and I will try for *Corel PhotoPaint*. All fine so far, but I look in the shops and where is *Encylox*, *Dangerous Tuxes*, *Linkxus*, *Family Trunx*, *Clipart for Penguins*, *Penguin Print Parade*, etc? So far Linux is something in our house that dad uses because he is a tinkerer, but I have a wife and two daughters who have their own real computer and all the programs money can buy.

And at work we run DOS programs through Win X/Y/or Z on a Novell network, as well as some Windows business apps, and there are very few crashes, our network administrator has just had holidays and there wasn't a problem on approx 50 PCs on the network. Just one video server running NT4 in a theatre presentation froze once.

It may be a 'free' system but I could buy the ME upgrade for the cost of the mags I've bought to get as far as I have with Linux.

"...you'll never go back to your old OS ..." to quote a sub-heading on the October cover, (it takes a while to get to Aus.) I still haven't left. But, I spend more of my time on a 'drive' far away.

Ian Caldwell
Victoria, Australia.

Well, maybe Linux isn't for everyone,

but can I ask what you actually want from Linux? Setting up a Linux distribution to work the way you want it to can still be hard work. Particularly with distributions like Corel, which are excellent, seamless and easy to install – providing you have the right hardware. Try to install it on a machine with an unusual SCSI interface or graphics card, and you'll get nowhere fast.

But rather than swap installations constantly, why not build on the ones you have? The software is pretty much the same (apart from some configuration utilities). For example, any of the distributions you mention are capable of browsing your Windows partition – you just have to mount it. Many installers will set this up for you, but you can still do it with Red Hat 6 or whatever.

This is one of the gaps in the Linux experience that this magazine tries to fill. As Linux becomes more popular, and manufacturers thus become keener to have their hardware supported (which is happening), this experience might change. But if you stick at trying to get a system working the way you want it, rather than just installing every new distribution that comes along, you might find things easier!

Easy install

I have been a regular reader of your magazine since issue two. It is well laid out and very informative but, I still have the feeling that something is missing that the other PC mags have. No not the overwhelming amount of advertising – I think you have done very well to keep this to a minimum. It must be the way you present the cover disc contents. I always look forward to anything free and a CD with various programs and so on is always of interest but, I find using the contents and your advice on using the contents frustrating. I have tried a few of the bits and pieces on several of the discs with varying degrees of success but, on the whole I feel let down by the experience.


I hate to compare it with other magazines since they are mostly concerned with another OS I am still

trying to get away from but, for now, they have the edge I'm sad to say. This is because for me, they explain how to install things and they usually work.

If I open the contents in one of the magazine's cover discs I can simply click on the utility or program that I want to install and away it goes or I have the option to unzip it to somewhere. If I do the same with the *Linux Format* cover disc I am often left sitting at a computer screen that doesn't do anything other than tell me what it is all about eg, your last inclusion of the 2.4 kernel. I was able to work out what to do with it from the tutorial in the July issue but, this should be much more automatic surely. Please could we have more automation in the installation of the files from your disc, it would go a long way to making life with Linux easier.

Thanks,
Derek Logan,
via email

I'm sorry you don't find using your Linux Format CD as fulfilling as it should be, but the reasons you mention are largely to do with the nature of Linux. As a multi-platform OS, most software is often provided as source code, which has to be compiled on the host system to run. Complicating this is the fact that virtually every Linux box is different to every other in some respect, and also that some software has specific installation requirements – e.g. config files to be edited and such. Even when the authors provide precompiled binaries in RPM/deb form, installation is not always as straightforward as you might like. Until there is a genuine, easy to use installation system widely adopted for Linux, there is unlikely to



The LXF CD is always packed to the gills with good stuff, but is it too difficult to get to?

be a 'click and install' option for anything on the CD anyway.

Hopefully you will find the instructions in our CD pages more useful. Again, unfortunately it is impossible to print sets of instructions for every Linux distribution, which is why whenever possible we always include source, as this will always work (provided your basic Linux installation supports the requirements of the software). We have made some efforts to make the CD pages easier to follow no matter what Linux system you are using, but the click and install interface is quite a way off for Linux.

Different strokes

As a regular reader of your magazine I feel as though it is time for me to write in and air one minor criticism about the magazine!

This criticism lies with the tutorials which you regularly run regarding installation and configuration of software. A prime example of this is in the January 2001 issue, in which there is an article on the *Roxen* web server. Whilst I was following this article it became apparent to me as a long time SuSE user that all of the 'init' scripts are called differently in SuSE than in other distros e.g. Redhat. My worry is that with even more Newbies turning to a version of SuSE as their first Linux distro, whether due to hardware support or ease of installation, I feel as though Newbies are probably feeling let down by not

being able to achieve the same results as you.

Therefore, may I suggest that you conduct a survey of readers in order to ascertain which is the favoured distro. You could then publish the results and tailor the tutorials accordingly. Hell, you could even turn it into a competition!

Mr Adam Parker
via email

This is a valid point, and one we try to get around by explaining major differences between distros. One of these is the way init scripts are set up, and, in the instance you mention, this is perhaps not clearly pointed out. Sadly there are so many variations of Linux that it's not always possible to include examples that will work on each setup, but I agree that we should highlight where this is the case.


Request

I'm learning Java but can't seem to get Kaffe working. So how about putting the JDK 1.3.0 and Forte for Linux on the next cover disc from the Sun website. It's a 36Mb download but my Surftime link won't stay up that long, pppd just dies!!

Mike Surtees

Sadly neither Forte or the JDK are GPL'd and, as such, require permission from Sun to be redistributed on our CD. Until they grant us a free license to do so (more likely than it used to be) your best bet is to get a decent ISP service that will stay running long enough, or buy a boxed set distribution which includes the JDK.

What is the problem with Kaffe anyway? I'm sure we can sort it out for you if you let us know...



We can't include Sun's JDK on the LXF CD... yet.

Eye strain

Thanks for a great magazine, but please avoid certain colour combinations if you really want us to read it all. Black text on a red or dark blue/green background is extremely unkind to the eyes and are sometimes unreadable. Examples: December 2000 edition, page 87 (Installing *SmoothWall*) and page 89 (small panel titled *Star Office 5.2 v OpenOffice*).

Your directions on installing and configuring *SmoothWall* were clear. But the article could have been richer if you had picked up some of the experiences posted in the SourceForge forum. For example, the modem attached to the *SmoothWall* box must be on the first serial port (COM1), otherwise it will not work.

Raph Awoseyin
rsawoseyin@usa.net

Point noted. I'll pass your colour clash queries on to our new art boy, who will no doubt pick out some more pleasing colour combinations.

Thanks for the SmoothWall tip.

School's out for Linux

I am the SysAdmin for a school in Devon. It is my heart's desire to migrate our system to Linux from Win 95. We currently use SuSE 7.2 on one of our servers, and on a couple of workstations. Now, I am not quite a Linux newbie any more, but there is a long way to go before I can fix problems that may arise the way I can with Windows.

I feel that we cannot really win over the next generation of users without being able to give them a (relatively) easy ride into Linux – the set-up and administration of which is difficult compared with the automated setup of Windows. The situation is changing; 'intelligent' installers are appearing with new distributions and are continually improving. SuSE 7.2 doesn't actually need any complicated questions answered for a 'Newbie' setup. Installing ReiserFS, however, needs expert mode – at the moment – maybe more newbie-help-wizards would help?)

Sadly, many people forget what it was like to be learning as a school pupil: the hours seem long, the tasks are tedious and interminable, and you have to remember so many pointless facts. So, who cares how you apply an SQL filter and sort on a new Table when you have to remember what Henry VIII's wives were called for a test next lesson?

Unfortunately, whatever anyone says about Microsoft's products, they have a relative simplicity for the young (and inexperienced) user. We, as a school, need to introduce pupils to the world of wordprocessors, spreadsheets and databases rapidly and painlessly – and the Office suite provides a common look/feel and scalable tools (Access Database, particularly – databases are difficult for pupils to grasp).

The two major issues which affect our migration are:

1. The (apparent) lack of support for the Linux platform from the makers of educational and reference software – it is difficult to find a CD-based encyclopedia which runs under Linux. Reference software is a major headache! Would HTML-based reference works be so awful? (No, is the simple answer). Any help in this direction would be much welcomed by us all.

2. The lack of a fairly simple, integrated, small-scale database system (tables, queries and forms in one place on the desktop) which is comprehensible by teachers, administrators and pupils alike.

The second issue seems to have been (partially) answered by *Star Office 5.2* with *Adabas D*, but the interface is still a little complicated for our users. It is, however, getting there. A simple database front-end, user administrable, is really what we need – spreadsheets, wordprocessors and graphics software there are aplenty.

I hope and suspect that by the summer of 2001 or 2002, (such things happen during the summer holidays), I will be in a position to tell my Head of Department that Linux is an educationally viable upgrade for our systems (as well as being cheap – a major consideration).

Already, at least one of our pupils has dumped Windows in favour of Linux, and others in the schools are dual-booting. It would be great to be able to give them all a headstart in that direction.

So I ask: Would *Linux Format* give us schools users who want to change a few pointers?

I hope the diatribe isn't too obscure or excessive. (Incidentally, has anyone come up with a decent random abuse or insult generator for Linux? – just asking...)

Regards,
Paul Ellison,
Edgehill IT

I can see your first point totally. There is very little in the way of educational software for Linux, and we are not aware of any developers working on major projects. There are things like Free Physics (freephysics.sourceforge.net), but they're few and far between.

In terms of encyclopedias, why not consider online ones such as the Encyclopedia Britannica, www.britannica.co.uk (which sadly has entries for Gates, William; Microsoft et al, but no entry for Linux) or Comptons Online Encyclopedia, www.comptons.com. I'm not sure of your Internet setup, but if you are permanently hooked up, these are pretty good options.

It may be a little beyond the scope of what you need, but it would be possible to create a customisable database interface (which you could then make as simple as you wanted) using an SQL server and a PHP web interface – though obviously this would require a fair degree of knowledge. There are some examples of this sort of work on Freshmeat.net, which may be adaptable to your purpose (there is certainly a simple address book written in PHP, called, if I remember correctly, Address Book).

Because of its affinity to servers and the Internet, Linux is a bit lacking in flat file database applications. I think the PHP/SQL route, while it may seem overcomplicated for your needs, might provide a more flexible solution.

Do any of our other scholastic readers have suggestions? **LXF**

FEATURE | 3D graphics



MAKING 3D GRAPHICS

work with Linux

So you've shelled out for a top of the range 3D card, but now what? **Linux Format** takes you through the basics...



Quake III Arena allows texture mapping onto realistic curved surfaces.

The biggest revolution in the computer games world was the introduction of 3D. Instead of a little sprite running from left to right on the screen, we now had shapely Lara Croft, terrifying Shamblers and slick Formula One cars. Graphical complexity bred better 3D cards (or was it the other way round?) which, in turn, led to a more immersive experience for gamers.

But it's not only trigger-happy gamers who have seen benefits in the rapid evolution of 3D technology. Off the shelf technology can now be used to model, animate and render entire films, creating sets – and even characters – that would otherwise cost millions to build.

The key to making the most of the capabilities of the latest generation of graphics cards – which can shift around millions of polygons per second – is ensuring your machine is set up correctly to use it. Although this is becoming easier as manufacturers discover the value of Linux, it's still not a trivial task.

One problem is that 3D hardware accelerators vary greatly in the features that they provide and the way that they are controlled. It would be ridiculous for each application wanting to use 3D

hardware to need rewriting for each accelerator. Fortunately, various libraries have been created to abstract away the details of the hardware and provide a common interface. The programmer then writes his 3D application to use this interface, and leaves it up to the hardware manufacturer or the library vendor to provide drivers for individual cards. Popular 3D libraries include Microsoft's Direct3D and OpenGL; the latter has become the standard on Linux and other Unix systems.

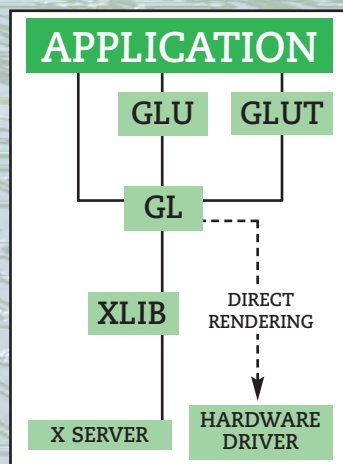
What is OpenGL?

The Open Graphics Library or OpenGL is a cross-platform interface that supports 3D rendering and hardware acceleration. It provides applications with a high performance toolkit that has functions for creating and manipulating 2D and 3D geometrical objects and rendering those objects to a frame buffer with advanced effects such as shading, texture mapping, fogging and motion-blurring.

OpenGL was initially developed by SGI from their earlier IRIX GL system and first introduced in 1992. Today OpenGL is virtually an industry standard with versions available for every major operating system, and the OpenGL

specification is now maintained by the OpenGL Architecture Review Board – whose members include SGI, Intel, IBM and Microsoft. SGI license the OpenGL technology to software and hardware vendors for implementation and provide standards validation on those implementations.

One of the key advantages of OpenGL is its portability; it is not tied to a particular operating system or windowing environment. The job of integrating OpenGL within a windowing system is taken care of by various extensions. One such extension, the GLX protocol, specifies how GL is married to X.



3DNow!


Starting with the K6-2 processor, AMD introduced a new technology in their processors called 3DNow! The aim of this was to provide more efficient processor functions to handle common floating point instructions often used by multimedia applications. It does this by providing 21 new instructions which can be used by any application.

So can 3DNow! make any difference to Linux? The short answer is: yes. Because the technology does not rely on anything outside of the processor itself, it is always on, and always working, but obviously programmers need to take notice of it. The kernel itself does not need to enable 3DNow!, although the 2.4 kernel does utilise some of the extra instructions in AMD's latest Athlon processors for some memory functions.

3DNow! makes more of a difference on Windows systems when it comes to specific 3D cards, where the drivers for the cards have been optimised to take advantage of the technology.

However, some individual applications do support 3DNow! for Linux, including the XMMS mp3 player and 3D applications such as Gizmo3D, a development tool for creating 3D software.

FEATURE | 3D graphics



→ GLX describes an extension to the X protocol which allows OpenGL commands to be sent to an X server and allows rendering into X windows or pixmaps. Since GL commands are sent over an X wire to the server, GLX is network transparent. That is, applications can redirect their output to be rendered on a remote machine. The GLX protocol also specifies a special case when the application is running on the same machine as the 3D hardware itself, and, in this case, can by-pass the X server. This is a technique known as Direct Rendering and allows the highest performance from the hardware.

SGI open-sourced the GLX protocol in February 1999 and the OpenGL system in January 2000, but they still maintain a reference implementation of both standards. This reference implementation doesn't support Direct Rendering or, indeed, hardware acceleration under Linux. That is the job of other projects.

Beside the core rendering API, OpenGL is usually complemented by some higher-level libraries as well, including GLU and GLUT. GLU, or the

GL Utility Library provides extra functions not strictly related to rendering such as image manipulation, matrix transformations, handling NURBS and so on. GLUT, or the GL Utility Toolkit, is a third-party library written by Mark Kilgard and provides an API for accessing windowing system functions that are independent of the particular windowing environment used.

There are three distinct solutions for driving 3D hardware under Linux

OpenGL and Linux

Unless you buy a commercial X server for Linux, such as X or Y, then you probably won't actually have OpenGL installed on your system. What you will have, however, is a version of Mesa. The Mesa 3D project is a third-party clone of OpenGL begun by Brian Paul in 1993 and initially released under the GNU LGPL. Mesa is not an official version of OpenGL since it has not been validated by SGI, but it is a very good substitute. The current release of

Mesa supports the OpenGL 1.2 specification, provides implementations of GLX and GLU, and supports GLUT.

Mesa started life with no support for 3D hardware acceleration and performed all of its rendering in software. More recently, however, versions have been created with drivers for certain hardware accelerators.

There are three distinct solutions for driving 3D hardware under Linux. The first is Mesa-Glide, which is a build of Mesa which supports 3dfx's Voodoo accelerators and makes use of the Linux port of the 3dfx Glide library. The second is Utah-GLX, a hardware-accelerated implementation of the GLX protocol for version 3.3.x of XFree86. Utah-GLX is shipped with drivers for a range of 3D chipsets and incorporates Mesa as its client interface. The third solution is DRI, or the Direct Rendering

Infrastructure. This project was designed to bring true Directing Rendering to Linux, and features a GLX implementation, Mesa and a Direct Rendering

Manger for the Linux kernel to synchronise access to the graphics hardware. DRI was incorporated into XFree86 with version 4.0.

Driving your hardware

The easiest way to get your 3D accelerator working under Linux is to get a distro such as SuSE 7.0 or Mandrake 7.2 which can detect 3D hardware and install and configure an appropriate driver. If you don't have such a distro, have a 3D card which they don't support, or you just want the satisfaction of having done it yourself, then follow our guide.

Mesa-Glide

The first 3D chipset to be supported with a version of Mesa was 3dfx's Voodoo range. This was made possible through a Linux port of 3dfx's Glide library, the proprietary API used to drive 3dfx cards. Two principal versions of the Glide library are available: Glide2 supports the Voodoo Graphics, Voodoo2, Voodoo Banshee and Voodoo3; Glide3 supports the Voodoo3, 4 and 5. The Glide libraries

SUPPORTED 3D HARDWARE

3D Chipset/card	Supported in Utah-GLX	Supported in DRI
3dfx Banshee/Voodoo3+	No	Yes
3dlabs Oxygen GMX/2000	No	Yes
ATI 3D Rage Pro (Mach64)	Yes	No (not complete)
ATI Rage 128	No	Yes
Intel i810	Yes	Yes
nVidia TNT	Yes (no AGP)	Yes (with drivers from nVidia)
S3 Savage3D	Yes	No
S3 ViRGE	Yes	No
SiS 6326	Yes	No
Sun Creator/Creator 3D	No	Yes

are available from <http://linux.3dfx.com/> or via CVS from <http://glide.sourceforge.net/>.

Mesa has included Glide drivers since release 2.2, and this offers one method of using a Voodoo card under Linux. However, unless you have a card not supported by Glide3, you'll probably be better off using DRI and XFree86 4.0. The Mesa-Glide drive is rather primitive by comparison.

To be able to use hardware acceleration with plain Mesa and a Voodoo card, you'll need a version of Mesa built with the Glide driver. If you cannot find a binary distribution of this, you'll have to build it yourself from the Mesa source code available from www.mesa3d.org. First, you'll need to get and install the Glide developer's kit from 3dfx, and, when you've done that, the Mesa configure script will detect it and automatically build Mesa for Glide. Glide provides no special access to Voodoo hardware. As it stands, clients of the Glide library on Linux need to be run as root to be able to access the hardware at all. To overcome this, there is a kernel module available which provides a driver to give non-root access via the device file `/dev/3dfx`. This is available from the same places as Glide, and you'll need to build it against your own particular kernel image.

For more information on this see the 3dfx HOWTO.

Utah GLX

The aim of the Utah GLX project is to provide hardware-accelerated Mesa drivers under XFree86 3.3.5 and 3.3.6 for various 3D chipsets. It does not support true Direct Rendering (although there is a hack for a kind of pseudo-Direct Rendering) and it won't work with XFree86 4.0. At the time of writing it supports ATI 3D Rage Pro, Intel i810, Matrox G200/G400, nVidia TNT series, S3 Savage 3D, S3 Virge and SiS 6326 chipsets. The ATI and Matrox drivers optionally support AGP, while AGP is required for the i810 driver. See the AGP boxout. It should be noted that the nVidia driver does not currently support AGP, so users with this type of card will find much better performance with the DRI drivers in XFree86 4.0. ➔

WHAT DOES A 3D CARD DO?


Rendering a 3D scene requires a series of relatively simple mathematical operations. Unfortunately, it requires an awful lot of them, and if you want to get a high animation speed, it is necessary to perform these many times a second.

The main CPU in your computer may be pretty fast, but ask it to do all this work and calculate a whole bunch of other stuff for the application/game you are running at the same time, and it will slow to an undignified crawl.

A 3D accelerator card basically contains a special processor, custom designed to process 3D operations. Because it has been designed specifically for this task, and has nothing else to think about, the 3D card can manage much greater framerates than the CPU alone, and even has some time to throw in some special lighting effects.

A frame from a game like *Quake* is first made up by creating polygon data for the objects in the scene. This polygon data is then passed over to the 3D card for processing. 3D cards from different manufacturers (i.e. different graphics processors), tend to treat this data in different ways, each trying to optimise the process in some way to make their processor the fastest. In general though, the next stage is that the processor determines which polygons are forward facing (i.e. ones which can be seen from the front of the screen) and discards the rest.

The card then retrieves 2D texture maps for the polygons and applies these to each of the polygons, while simultaneously deciding which polygons are 'in front' of others, using the z-buffer. The texture maps themselves may be fetched once and stored in memory local to the 3D processor. Because the polygons may either be close up or far away from the viewer, they are usually stored together as a set of textures of increasing levels of detail: a MIP map. The actual value used for each element in the texture map (texel) is then calculated by averaging values from between the different versions (called bilinear or trilinear filtering).




Lighting effects can add depth to a scene. Here we're navigating by plasma gun light.

At this stage the polygons are all where they should be and textured properly, but still look quite flat. The next stage is to add lighting effects to bring a degree of depth to the objects. A common way to do this is to use Gouraud shading, a very fast algorithm which can be used to change the 'brightness' of each element on the polygon. This system calculates the shading levels at all the edges of each polygon, then interpolates the results across it. At this stage bump maps can be applied to introduce surface irregularities and make the polygons more realistic and less like flat planes. Any special lighting effects (which vary greatly between processors) like fog and translucency are usually applied at this time too.

The final stage is to combine the texture, lighting and colour values together to create a single pixel. A lot of 3D hardware does this at a slightly higher resolution than the intended display, and scales the result down to remove or disguise small imperfections in the rendering process (which can occur where polygons intersect, or share vertices) and to antialias the final image.

The image, now ready for display, is often put into a frame buffer. The buffer allows one or more completed frames to be queued up for display. The card then switches from displaying one




It looks so simple, but your average 3D card has a lot of work cut out for it when you're playing Quake.

image to the next, freeing up the old frame memory to be used for another buffer. This technique is called double-buffering, and is used to make the animation of frames much smoother (otherwise you'd be able to perceive the frame being created).



Texture maps give walls, drinks machines and monsters a more realistic sheen.

FEATURE 3D graphics



The nVidia TNT series of cards are great, but you'll need nVidia's own drivers to run them.

→ Utah-GLX consists of two parts: a specially-built version of Mesa and a GLX module for the X-server. You need to obtain or build both for a working installation. The Utah GLX source code is available via CVS at <http://utah-glx.sourceforge.net/>. This is probably the best route to go, since the project hasn't achieved a stable release yet. You can also get source code snapshots from <http://snow.ashlu.bc.ca/glx/snapshots/>, and some Linux distro vendors make recent builds available as binary packages.

To get the code by CVS, log into the CVS server with:

```
cv$ -d :pserver:anonymous@cv$
.utah-glx.sourceforge.net:/cv$root
/utah-glx login
```

When prompted for a password, just hit Return. Next, download the GLX source tree with:

```
cv$ -d :pserver:anonymous@cv$
.utah-glx.sourceforge.net:/
cv$root/utah-glx co glx
```

and download the Mesa tree with:

```
cv$ -z3 -d:pserver:anonymous
@cv$.mesa3d.sourceforge.net:/cv$
root/mesa3d co -r mesa_3_2_dev
Mesa
```

This will create the directories 'glx' and 'Mesa' in your current directory. Link the Mesa tree to the GLX tree with:

```
ln -s Mesa glx/mesa
```

and you are ready to build. First, configure with:

```
cd glx
./autogen.sh --enable-extra
```

By default, Utah-GLX libraries will be installed to `/usr/local/lib`. If you want to change this, check the configure options with:

```
./configure --help
```

and call **configure** again. Typically, you will want the client libraries installed over your current GL libraries (either `/usr/lib` or `/usr/X11R6/lib`) and the X server module in `/usr/X11R6/lib/modules`. So, for example, enter:

```
./configure --enable-extra --
prefix=/usr --with-
moduledir=/usr/X11R6/modules
```

Now build and install Utah-GLX:

```
make depend
make
make install
```

You'll now need to modify your X configuration to load the GLX module. So add the line **Load "glx.so"** to the "Module" section of `/etc/X11/XF86Config`. If you have installed this module to a non-standard place, you will have to supply the full path in the above line. You may have also to modify some of the options in the "Device" section, depending on which 3D chipset you have. Consult the

Utah documentation for more details. The Utah-GLX settings are stored in `/etc/X11/glx.conf` and may also have to be adapted for your set up.

Restart your X server and you'll be ready for hardware-accelerated 3D.

DRI

DRI stands for the Direct Rendering Infrastructure and is a project initiated by Precision Insight to create full OpenGL Direct Rendering for XFree86. The first release of XFree86 to feature DRI was version 4.0.

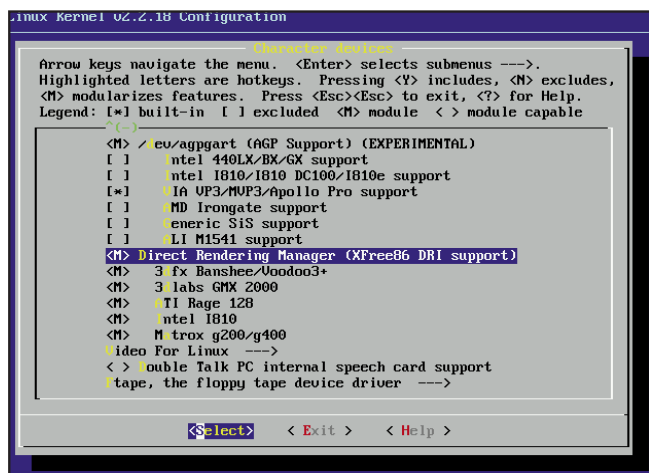
DRI consists of several separate pieces: the GLX extension itself (implemented as a module for the X server), the 3D drivers for the various chipsets supported (the correct driver is dynamically loaded by the GLX module), the client-side API (yet another version of Mesa), and the Direct Rendering Manager (DRM). The first three components are available as part of the XFree86 4.0 distribution while the DRM is a kernel module and comes as standard with 2.2.18 and 2.4.0 kernels.

For Direct Rendering to work, it is essential to have an appropriate Direct Rendering Manager built into the kernel. To enable DRM support in the kernel, turn on `CONFIG_DRM` from the character devices page and select the DRM for your graphics architecture. The Rage 128, i810 and Matrox drivers all require AGP support in the kernel.

Currently the DRI project has DRI and DRM drivers for 3dfx Voodoo Banshee3+ (this requires the Glide3 library), 3dlabs GMX 2000, ATI Rage 128, Intel i810, Matrox G200/G400 cards on x86 platforms and for the Sun Creator 3D on the Sparc. DRI is currently being ported to the Alpha architecture and drivers are promised for the ATI 3D Rage Pro and the ATI Radeon. In addition, the nVidia TNT range is supported by proprietary drivers by from nVidia themselves.

If you have a correctly configured kernel and a properly installed build of XFree86 4.0 or later, then the drivers are easy to set up. All you need do is make sure that the lines

```
Load "glx"
Load "dri"
```



Direct Rendering under Xfree4 requires a rendering manager for your graphics hardware to be built into the Linux kernel.

HOW FAST DOES IT GO?

When dealing with animated 3D applications (mostly games, but also some useful stuff like CAD software), you want to be able to achieve a frame rate of at least 20fps (frames per second) to make the viewer perceive the action as realistic. 50fps gives a very smooth animation, and although faster speeds are possible, a frame rate over 80fps isn't really any better, because that's around the limit at which any differences are perceptible.

The achieved frame rate obviously depends on a number of factors, including the complexity of the image being rendered and the speed of the CPU and graphics card. The CPU speed still comes into the equation, because the main processor still has to calculate and provide the 3D card with instructions on what to draw. As 3D processor technology has increased, the bottleneck no longer appears in the ability of the 3D card to process the data, but the speed at which the data can be delivered to the processor.

This is the reason behind the

development of the Accelerated Graphics Port (AGP) bus in modern computer hardware. AGP was developed by Intel to provide a high bandwidth way for the graphics processor to access data stored in the main system RAM (large textures etc)

quickly. Most new motherboards offer AGPx2, which means they can transfer data on both the rising and falling edges of the bus clock signal. On a standard 66MHz bus this allows a throughput of 528Mb/second. The 3D subsystem is still reliant to some extent on the speed that the actual data of what to draw can be supplied to it though, which relies on the main CPU. Although a 3D accelerator makes a huge difference, the CPU still has to do some of the work!

If you can't achieve a usable



Games become more sophisticated with every passing month.

framerate, you can try a few things to reduce the complexity of the image being displayed (before you shell out on a faster CPU or new graphics card). Usually the software itself will allow you to reduce the screensize, and the complexity. Turning off special lighting effects will increase the speed, as will reducing the bit depth of the screen (although some new cards work internally at a fixed depth, so this may not have any noticeable effect).


are included in the "Module" section of your `/etc/XF86Config-4` file and everything should work automatically. Consult the DRI documentation for any details specific to your graphics chipset.

The nVidia drivers are typically more complex to set up. The latest nVidia kernel and XFree86 4.0 modules are available from ftp://ftp1.demon.net/nvidia.com/pub/driver/s/english/XFree86_40/0.9-5 RPMs are supplied for a range of standard kernel images from a selection of distros, but if you are not using one of these, then you'll have to get the `NVIDIA_kernel-0.9-5.tar.gz` archive and build the module against your kernel image. Unpack the archive and enter **make** to build – or at least glue – and install the module. Note that this module may cause problems on 2.4 series kernels.

The nVidia XFree86 modules are also available as an RPM or as a binary tarball. To install the latter download the archive, unpack it and enter **make** to install it. This will copy the various drivers to the correct places, including

the GLX extension module, the 2D driver and the 3D driver.

There are some points to be aware of. Firstly, the nVidia driver kit replaces the standard XFree86 4.0 `nv_drv.o` 2D driver with a driver called `nvidia_drv.o`. You might want to remove the old driver (in `/usr/X11R6/lib/modules`

`/drivers`) just to make sure you don't get confused, and remember to change your `XF86Config-4` file if you have explicitly named the driver to use. Secondly, having the old GLX module hanging around can cause problems too. So, remove the file `/usr/X11R6/lib/modules/extensions/libglx.a`. 



It's not just shoot 'em ups that benefit from 3D acceleration.

AGP

AGP or Accelerated Graphics Port is a bus specification by Intel which allows graphics cards faster access to the system's main memory than the standard PCI bus. The AGP controller acts as a kind of memory management unit for the graphics card, making discontinuous blocks of system memory appear as a single chunk, and enables burst access to this memory. The AGP runs at twice the PCI bus speed (66 MHz), with AGP 2x doubling this up to 133 MHz.

The Linux 2.2.18 and 2.4.0 kernels provides a `/dev/agpgart` driver (AGPGART stands for AGP Graphics Aperture Relocation Table) which a particular GLX driver can use to send commands to the card. To include AGP support in your kernel, enable the `CONFIG_AGP` option (in the character devices section) and enable support for your motherboard chipset. Chipsets covered include most Intel chipsets, the ALi M1541 and VIA chipsets like the Apollo Pro.

You will also want to enable the `CONFIG_MTRR` or Memory Type Range Registers option as well. The MTRR registers on the Pentium Pro and later CPUs (non-Intel CPUs have equivalent registers) can be used to control processor access to memory ranges. Your 3D driver will use this driver to enable write-combining and combine small memory transfers into one larger transfer and so boost performance.



No matter how much you think you know about Linux, there's always something more to learn. This is a good thing because a) it gives us plenty to write about and b) gives you the opportunity to try out new things all the time. How cool is that? This month we have a detailed look at the customisation capabilities of sawfish (the Window manager formerly known as sawmill), uncover the power of regular expressions in Perl, multiple threads in Java and more.

Richard Drummond

TUTORIALS

You'll never tire of finding out what Linux has to offer you, and we'll never tire of giving you our easy-to-follow tutorials. **LINUX Format** will turn you into a Linux guru in no time!

IN THIS ISSUE...

WINDOW DECORATION

Sawfish tamed

Whether you want simple windows, or an all-singing resource hungry fantasy world on your desktop, Sawfish could be the window manager for you. Biagio Lucini shows how to get the best from it...

p68

Don't get stuck with a boring desktop, make it pretty!




TEXT PROCESSING BASICS

Perl explained

Charlie Stross continues his look at Perl with an examination of the way the language handles text processing tasks such as find/replace using regular expressions, and the use and enforcement of namespaces.

p72

Don't let Perl's word handling give you the hump (ha ha!)



LINUX FORMAT

Let us know

If there is anything in particular that you would like to see a tutorial on, the best thing to do is just ask. There are literally thousands of things we could do tutorials on, but we want to give you the ones you most want to read. Whatever the topic, if there is enough demand, we'll run it. To make your request, please contact us here at:

LINUX Format
Future Publishing
30 Monmouth Street
Bath BA1 2BW

email:
linuxformat@futurenet.co.uk

JAVA PROGRAMMING

Threads: Part 1

Java is great at handling more than one job at a time. In part six of his tutorial, Richard Drummond begins to delve into the world of threads to discover how Java can slice up time to create a very good illusion of

p76

Don't just do one thing at a time with Java.

```
bash-2.04$ java ThreadDemo
Thread-0 is running
main is done.
Thread-1 is running
Thread-1 is running
Thread-1 is running
Thread-1 is running
Thread-0 is running
Thread-0 is running
Thread-1 is running
```



COVERDISC QUIRKS

Installation support

Over the past ten issues we've given you a number of popular distros to choose from. Simon N. Goodwin uncovers the frustrating quirks that may have put you off trying Mandrake, Corel, Storm etc, and offers solutions to some of the common problems.

p79

Don't let Corel's installer give you nightmares.



RATINGS BOXES EXPLAINED

Some people like to know how tricky things are going to be before they start, so we have decided to

split our tutorials into three categories, based on your familiarity with Linux.

Please don't let this put you off reading anything – they're just for your guidance!

DIFFICULTY LEVEL

Fresh-faced Newbie



DIFFICULTY LEVEL

Stubble-faced Geek



DIFFICULTY LEVEL

Full-bearded Hacker



What on Earth is LDAP?

LDAP stands for “Lightweight Directory Access Protocol” and is pronounced ‘el-dap’. Dave Coulson reveals the hidden depths of this (sort of) hierarchical database.

LDAP is a protocol for accessing directory based information over a network. Information on an LDAP server is stored with a hierarchical structure, rather than a flat ‘column and row’ format like a normal database, which lends itself to storing large amounts of similar details in an organised fashion. Originally, an LDAP server was a thing which sat between an LDAP client and a DAP (X.500) server, so as to reduce the resources required to run the client. As LDAP provided fewer features than DAP, people eventually gave up using DAP for all but the most powerful applications, and just ran an LDAP server on it's own to handle everything. LDAP, being directory orientated, is first used to find the information you want – letting you know where it is in the directory structure – and then you tell it go and fetch the

data. One of the main advantages of this is that you can transparently distribute this over many machines, so unlike traditional load balancing techniques, you can have separate parts of the directory on different machines. It is, in a way, analogous to storing different directories in the Linux root filesystem on NFS mounts from different file servers. This gives scalability which many other data storage systems do not offer as you can scale, either up or down, individual components within the structure without necessarily changing anything else.


SO, IT'S LIKE A DATABASE?

You can think of it like a N dimensional database, but it's a lot easier to think of it as an tree style structure, with one object being the child of another, and so on. If, for example, you were storing personnel records, it would have the company at the top, then each individual department, then managers and would continue down the hierarchy of power in the organisation until it reached someone who was not senior to anyone else. Storing an object for each person isn't much use unless you actually give it some information about them, so we can set variables within objects to store specific information, such as their name, phone extension and such like. In fact, each person who is a ‘child’ to their initial object is actually the same as any variables stored within the person's object, as a variable is simply another branch in the tree.

CAN I STORE ANYTHING OTHER THAN PEOPLE IN IT?

Of course you can, as long as it's a structured hierarchy, much like a filesystem on a computer. You might store birthdays within it, but due to the ‘flat’ nature of the information it would be far more efficient to use a proper database. If you were an on-line retailer you may

LDAP stores your information in a tree-like hierarchy.



choose to store your product lines using LDAP, so you can organise each component into easy to access groups, which would be reflected in the front end of the service.

Within each object referring to a product you might store the manufacturer, a reference, price, a short description and anything else pertaining to that specific product type. You might want to get a list of all products from a specific manufacturer, with the list of contact details and a reference number within a separate area of the directory, which can be done quickly by telling the LDAP server to 'return all objects with manufacturer="RedHat" and are below such and such an object'. Even if you drift ten levels into the directory with products, you can quickly get the 'path' to an individual object quickly and easily.

HOW IS THE DATA ORGANISED?

Every item of data with LDAP is organised in a tree style structure, with each item, or entry, having it's own 'Distinguishing Name', or 'dn'. Every entry has a type, typically known as 'objectClass', set but every other attribute to an entry is simply a sub-entry to it. Unusually there is a single DN which everything is a sub-entry for, although if many people are sharing an LDAP server and the data needs to remain totally separate, there may be multiple top level DNs. So, we have a top level object with the dn 'o=ixf', the 'o' being totally arbitrary and meaning 'organisation'. We could, of course call it 'organisation=ixf', but if you use long DN variables and work your way through a good few levels, the DN for lower objects becomes quickly unmanageable. Under 'o=ixf', we might create a list of issues, the first being 'i=1'. So, we have 'i=1' as a sub-entry of 'o=ixf', so the final DN for the first issue is 'i=1, o=ixf'. DNs are written with the highest level object on the right and a comma separating the individual entries. If we stored an article within an issue, with the variable 'a', we could then have the DN 'a=anarticle, i=1, o=ixf'. Of course, we can have different types of entries as children of another, so we may have a 'c' variable for each issue for everyone who contributed to that specific issue.

As mentioned earlier, each entry has an **objectclass** variable, which distinguishes that specific entry's object type from everything else. We might have 'magazine', 'article', 'image', 'person', and such like; they're usually not very specific. Every entry will have its objectclass set as one from the predefined list, so if we wanted to search through all 'person' entries within the whole 'o=ixf' directory, it can be done quickly and easily. From a users point of view, 'objectclass' looks just like another sub-entry, but each entry can only have one, and it effects the functionality of the specific entry, such as the sub-entries that object is allowed.

HOW DO I SETUP ACCESS PRIVILEGES FOR THE LDAP SERVER?

Permissions for *slapd*, the LDAP server from *OpenLDAP*, are set within the `etc/openldap/slapd.conf` configuration file using the options **rootdn** and **rootpw**. At a first glance, you may assume that the **rootdn** applies to a real DN for a specific entry, but it is purely arbitrary and does not need to exist within the directory server. Keeping to our previous 'o=ixf' directory we might use the following:



```
rootdn "cn=root, o=ixf"
rootpw linuxformat
```

Then, after restarting LDAP, anyone who connects to the LDAP server will do so with 'root' permissions which, just like root on a Linux machine, means they can do anything they want. While this is great for the LDAP administrator, it's not much use for users, particularly if many groups are sharing the server, but we need it to remain secure. Each user can have their own DN, which has individual permissions set within the LDAP server. We could have a user 'cn=someone, o=ixf' and only have them access a specific directory DN, or type of entry. Firstly, we need to allow someone authorised by an entry to change the password for that entry. That is, to allow someone to change their password on the LDAP server.

```
access to attr=userPassword
by self write
by anonymous auth
by * none
```

Simply put, it will allow a user to change their own password, otherwise it will ask for authentication for anonymous connections. If you're logged in, but try to change 'userPassword' for someone else, it will fail. We can, if we want, add extra entries to that set so a specific user, such as an administrator, can change passwords. Adding:

```
by "cn=admin, o=ixf" write
```

would permit such a function. As always, after changing the contents of the `slapd.conf` file, restart *slapd* so it reloads the configuration. Even if a user forgets their password, you can still jump in with the **rootdn** and password to change it, so you never end up with the nasty situation where you need to reinstall it if you forget the root password.


LDAP, being a service, has the potential to be exploited, as with anything else, so as well as authentication restrictions, care should also be taken as to which user owns the *slapd* processes and which machines can access the ports, 389 with both tcp and udp, with the use of either `ipchains` or `iptables`. Using ➔

Any storage retrieval device needs a comprehensive indexing system.

What can I use to access the LDAP server?

Almost all popular, and not so popular, languages have support in one way or another for LDAP, including Perl, PHP and obviously C. The command line tools are such that it's fairly easy to use a shell script to do simple LDAP queries. You do need to have an LDAP distribution installed, so the headers are available for the program to build against, but other than that, it is fairly straight forward. PHP has support included in the main distribution, so you just pass /—with-ldap=/usr/local/openldap/ to configure and off you go. There are numerous LDAP modules for Perl, and a look on CPAN will produce a fairly comprehensive list.

WHAT ON EARTH | ldap



LDAP is ideal for storing and retrieving information about groups of people i.e your company's staff database.

→ tcp wrappers achieves a similar end, but leaves udp open for attack. Chances are that if you're using LDAP seriously, it will run purely on a backend server without any direct Internet access, so this should not be a problem.

HOW WOULD I USE LDAP FROM PHP?

Assuming you have already have PHP configured correctly with *Apache* and have built it with LDAP support, you can go ahead and start using it. First of all, though, we need to connect to the LDAP server using the **ldap_connect()** function, using the following method:

```
$ds=ldap_connect("localhost","389");
```

This will create a TCP connection, but not actually authenticate ourselves to use the server. In order to perform operations as a specific user, we need to tell LDAP who we are and what our password is using **ldap_bind()**

```
$r=ldap_bind($ds,"cn=root,dc=themes,dc=org","themes_org");
```

\$r isn't really of any use, but **\$ds** is the variable reference to the LDAP connection, which we need to specify when we search the directory. If we wanted to search our issue directory for an article titled 'ldap' we could search 'o=ixf' for 'ti=ldap' using **ldap_search()**:

```
$sr=ldap_search($ds,"o=ixf","ti=ldap");
```

This would return all entries under 'o=ixf' which have a sub-entry of 'ti: ldap'. Using **ldap_get_entries()** we can convert this into something more useful before stepping through the returned entries:

```
$info = ldap_get_entries($ds, $sr);
echo "Data for ".$info["count"]." items returned:<p>";
for ($i=0; $i<$info["count"]; $i++) {
    echo "dn is: ". $info[$i]["dn"] . "<br>";
    echo "first ti entry is: ". $info[$i]["ti"][0] . "<p>";
    echo "first c entry is: ". $info[$i]["c"][0] . "<p>";
}
```

If we had multiple 'c' entries, such as if there were multiple contributors to an article, this would only return the first one, but it is very easy to create code which does exactly what you want. Once we're done with our LDAP connection we can either ignore it and let PHP close it, or we can be a little neater and close it ourselves:

```
ldap_close($ds);
```

The PHP manual at www.php.net/manual has all of the LDAP commands for PHP, as well as excellent documentation of their usage.

HOW DO I ACCESS THE LDAP SERVER FROM THE SHELL?

OpenLDAP offers a number of command line tools to perform LDAP functions, either from a shell directly or via the LDIF data format. The simplest of these is **ldapadd** which does nothing more than add data to an LDAP server. It is invoked with a number of command line options to provide authentication, but after that it will just wait for data entry on stdin, unless you've told it to extract data from a file:

```
ldapadd -f entries.ldif -x -D "cn=Admin, o=ixf" -w password
```

or

```
ldapadd -x -D "cn=Admin, o=ixf" -w password
```

With the latter, it will not do anything to the data until **Control-D** is pressed, but the use of the LDIF format is useful if you're not totally sure what you're doing.

Searching is almost as easy, using **ldapssearch**. If we wanted to return the 'cn' entry for everything under the 'o=ixf' directory we would do;

```
ldapssearch -L -b "o=ixf" "cn=*"
```

Here, **cn=*** is a match for the entries we want to match and we can specify as many of them as we like, with LDAP matching them all. Doing:

```
ldapssearch -L -b "o=ixf" "cn=*foo*" co
```

would return all cn entries containing the substring 'foo' and the **co** sub-entry from each of those entries. As long as we pass the **-L** option, the returned information is displayed in the LDIF format, making for easy interpreting by a script or simply for importing into LDAP later.

CAN I INDEX CERTAIN ENTRIES?

As with any large scale database, indexing commonly used items is a must when there is a huge amount of data to search through if you don't want the server falling over when you try to access it. LDAP supports indexing on all types of entry, although if you index a string it's going to take more memory than an integer, in five different ways, so depending upon the type of search you will be performing, it may be useful to use a different type of index. The simplest of the indices is the **present** index, which is simply a boolean index of the existing entries, and as such, this is usually a default index for all data types. After **present** comes the **equality** index, which as you might expect, is a check for equality, so if you have a user database of 100,000 people, each with a uid, it's probably a good idea to index 'uid' for equality so searching for a specific person by uid is quick and efficient. For strings we have a **substring** index, which will speed up non-equality searches no end, however the memory costs are significant compared to numerical indices; so until it is really needed it may be best to steer clear, unless you have the necessary hardware.


The last index type, apart from **none**, is **approx** which will create an approximate index for the entries. As with a substring, it's not particularly resource friendly, but if you were searching a range of numerical values, it is useful. As with all the other server configuration options, these are setup in the slapd.conf file:

```
index default pres,eq
index objectClass,i
index o,co eq,sub,approx
```

Here, the default index is set to present and equality, and we have enforced that with the **objectClass** and **i** entry variables. For **o** and **co** we have set equality, substring and approximate, although the present index will be removed for these entries.

HOW DO I RUN IT ACROSS MULTIPLE MACHINES?

There are two ways to run *slapd* over multiple machines; one by replication, where load balancing is used to distribute the queries over different machines, the other by distributing a specified tree to a different server.



Starting with the latter, we create an entry with the objectclass **referral** and **extensibleObject** and a **ref** entry.

```
dn: dc=subtree,o=lx
objectClass: referral
objectClass: extensibleObject
dc: subtree
ref: ldap://ldap.box.net/dc=subtree,o=lx/
```

Easy, isn't it? As long as the other ldap server has the 'dc=subtree,o=lx' directory, distribution is totally transparent to the client.

Replication of *slapd* is done using *slurpd*, which distributes information from a master LDAP server to a number of slaves, so users can access the slaves without killing the master quickly. Firstly, we need to tell our master server to dump a load of all operations it does, so the slaves can do exactly the same. By setting the **replica** and **repllogfile <filename>** options in slapd.conf, anything it does will appear in a log file which *slurpd* can use to update the slaves.

The **replica** directive is fairly complex, as it is used to tell slurpd which slaves are updated with which directories, and could be setup as:

```
replica host=slave.example.com:389
binddn="cn=Replicator,o=lx"
bindmethod=simple credentials=secret
```

Here, the **binddn** is the username used for authentication and must exist on the slave LDAP server if it is not the same as the rootdn for the slave. The **binddn** variable is important, as within the configuration for the slave, any write functions done within the **updatedn** scope, which should be the same as the **binddn** will be sent to the master for processing, rather than going to just the single slave. Once this is done, we can shut down the master *slapd*, copy over the directory defined in slapd.conf to the slave, so they both start with the same thing. Then, we start the master and slave *slapd*, although at this point they will not be replicated automatically. If you perform a write option on the master server then check the replication logs, you will be able to see what it did in LDIF format. Once you run *slurpd*, the slave will become synchronised and everything will be happy. It is preferable to run the slurpd server on the same machine as the master server, so it doesn't need to use something extra such as NFS to give another machine access to the replication log. **LXF**

WHERE CAN I GET LDAP FOR LINUX?

The most popular LDAP distribution for Linux is OpenLDAP (www.openldap.org). This contains the server, *slapd*, as well as the replication server, *slurpd*. Installation is done with a slightly different `./configure &&, make depend &&, make &&, make install` routine and there are many options which can be passed to the configuration script, such as where it is to be installed and any specific options you want. There are, naturally, numerous commercial servers for Linux and other platforms, but if you're not sure if you need one of those, you probably don't.

A server isn't much use without a client, so there are also numerous client applications for writing, searching and reading on the directory server. One of the major things OpenLDAP has is the LDAP Data Interchange Format (LDIF), for storing directory entries within plain text files. You simply open a text editor, put in the dn of the entry you want to manage, then add, remove or modify any sub-entries before pushing the text file through *ldapmodify* to make the changes on the server.

Customising SAWFISH

Sawfish is a highly configurable window manager, thanks to its modular structure. **Biagio Lucini** reveals how to add some personality to your desktop.

We keep saying 'Linux is about choice' and this choice extends from applications, distributions and even desktops that can be configured for your way of working. There are also plenty of window managers around, each one trying to address a specific issue: some of them aim at a beautiful appearance (*Enlightenment*), some others try to recreate the interface of popular operating systems (this is the case of *WindowMaker*), others are inclined to being light-weight and fast (*Fvwm*, *BlackBox*).

The underlying philosophy of *Sawfish* can be summarised in two words: customisation and speed. In fact, the aim of this window manager is to allow for complete customisation, while remaining reasonably fast. Both these goals are pursued with a modular structure; every feature in *Sawfish* is a separate piece of code that the window manager executes only if requested. In this way, everyone can just keep the features they want, without loading unwanted bits. The result is a low overhead in terms of RAM, compared to other window managers (provided that not many extensions are loaded). In this 'search for speed' something must be sacrificed, but this is not a problem, for two reasons: firstly, a missing feature can always be implemented later, and secondly because the non-implemented features can always be performed by means of other utilities, remember that *Sawfish* is the official window manager of *GNOME*, and a lot of desktop settings can be achieved via this Desktop Manager. For these reasons, it's better to think of *Sawfish* as a way of removing the toys that

tend to zap resources rather than having a lack of features.

GETTING STARTED

Every distribution has its own preferred desktop. If you're using *GNOME*, it is quite likely that you are using *Sawfish* already. But if not, how can you start *Sawfish*?

There are several ways. Here we are going to describe the one that should work for everybody, even if this is not the most direct or the simplest way to get *Sawfish* working. First of all, you have to check if *Sawfish* is installed on your system. This can be achieved with the command **sawfish --help**. If you get a help message, you have *Sawfish*, but if you get the error 'bash: sawfish: command not found', you'll have to install it. Unless you have an old distribution, you will probably find the *Sawfish* package on the CD-ROM or on the ftp site of your distribution. If not, you can find it – with accompanying documentation – at the official home page of the project <http://sawmill.sourceforge.net> (the latest official release is 0.33.1, but this project is in continuous evolution). Once the above help command gives you a sensible answer, you are ready to proceed.

We'll assume that your Linux box boots in console mode. If you start X by default, then as root, open a console and type **init 3** (note that this works on RedHat and Mandrake, while on other distributions you have to check which runlevel corresponds to the multi-user text mode). You should now be prompted to login. After you've entered the system, edit your *.xinitrc* file (you should have it in your home directory; if not, just create one) and write in it just the line **sawfish** or **gnome-session** depending on whether you want to have


A background image can be set in a terminal using a dedicated utility.



SAWFISH OR SAWMILL?

"Sawfish" has not always been the name of the Window Manager we are describing: until a few months ago, it was called Sawmill. Then, the author discovered that a company was about to license a piece of software also called Sawmill, and decided to change the name. A big consequence was that in all the configuration files and directories the word sawmill had to be replaced by sawfish.

For instance, the personalisation file *.sawmillrc* was replaced by *.sawfishrc*. However, the new versions keep back compatibility with the old name, so they will work for instance with a *.sawmillrc* as well as with a *.sawfishrc*. The process of changing the name is not yet completed, so do not be concerned if you spot Sawmill where you were expecting Sawfish in some documents or Internet addresses.



The default appearance of Sawfish is a good starting point for the customisation of the window manager.

Sawfish working alone or embedded in the *GNOME* Desktop environment (if there are other lines in your `.xinitrc`, you can just comment them with a `#` character at the beginning of each line).

Save the file, then type **startx** and you're away. If you use *Sawfish* without *GNOME* (this is the suggested choice if your computer does not have at least 64Mb of RAM), you get a gray and white background, without any possibility of changing it from the Window Manager. This is just the default background of X. In other words, there isn't any background image because *Sawfish* does not set one up by default, however, this can be achieved by external applications. If you do not like such a plain desktop, you can always use *Sawfish* in conjunction with *GNOME*: just restart X with *GNOME*, then choose *Sawfish* in the window manager selection box. However, a pretty desktop has a price in terms of RAM usage, and can be unusable on older systems.

MODULARITY

There are two main characteristics that define a window manager: the possibility of manipulating windows and their appearance. Strictly speaking, the first is the main function of a window manager, but we can't ignore the desire for users to create their own special look and feel on the desktop. In *Sawfish*, the management of windows and appearance is configured using extensions to the basic application. These extensions are written in Librep, which is a dialect of the Lisp programming language. Librep is a programming language quite similar to Emacs, the language in which Emacs extensions are written in. If you followed our description of Emacs functions in Issue 7, you'll be familiar with this idea of modularity, which can be successfully achieved with a Lisp-like language. The main characteristic of these languages is that there is no real distinction between programs and data, since they are treated in the same way. This provides great flexibility: if you don't like *Sawfish* because it misses feature X, you can always add your own bit of Librep code that implements the feature you need. Indeed this is the way in which *Sawfish* works: it's a very basic window manager, extended with a lot of functions written in Librep. The *GNOME* compatibility is also achieved through a Librep module.

The main repository of Librep functions is (in Mandrake

7.2 at least) `/usr/share/sawfish/0.30.3/lisp`. This directory contains files with the suffix `.jl` and files with the suffix `.jlc`. The former contain the Librep source code, while the latter are the compiled versions that they can be executed rather than interpreted, translating into faster response times for. However, in normal operation, compilation of Librep code is not needed. In that directory, among other files there is `sawmill.jl`, which executes the loading all the required functions.

Hence, a way to add the new function `myfunction.jl` is to place the corresponding source code in the `/usr/share/sawfish/0.30.3/lisp` directory and then add to `sawmill.jl` the line (**require 'myfunction'**). Note that there isn't a second `'`. This procedure requires root privileges, and automatically makes the extension available to every user.

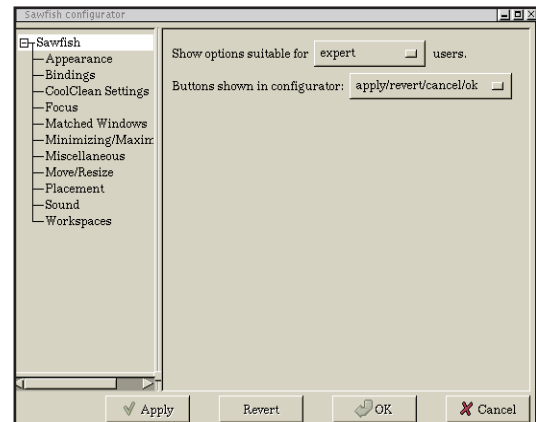
Below is an example of simple Librep code, however it is worth knowing that there is a repository of Librep functions for extending *Sawfish* at <http://www.sics.se/~lofgren/sawmill/>, so check that out before you spend a lot of time duplicating a function.

There is also a way to have an extension for private use. *Sawfish* stores the personal preferences of users in a directory named

`.sawfishrc` located in the home directory of each user. Inside here, there is (or can be created) a directory named `/lisp/`, where we can put our private Librep extensions. For instance in the previous example, if we do not have root privileges or we don't want `myfunction.jl` made available to the other users, we can put the file in `~/sawfish/lisp/` and then add the line (**require 'myfunction'**) to the file `.sawfishrc` (or `.sawmillrc` for an older installation). As for other files and directories mentioned here, if they don't exist on your system, they can be safely created. Now when you log back in, this new function will be available only to you.

THE CUSTOMISATION FILES

In the previous section, we have mentioned the file `.sawfishrc`. This is the main file in which the user can put their own customisation options. By default this file is not present, so if we plan to customise *Sawfish* to any degree, the first thing we should do is to create it. There is another file that is loaded only if the user does not have a `.sawfishrc` file: `sawmill-defaults.jl`. It is located in the usual Librep repository. Once a customisation file is created, it is a good idea to



Sawfish-ui can display customisation options for beginner, intermediate and expert users.



The infos are the main source of documentation about Sawfish.

Sawfish client

This program allows you to connect to a window manager session and to evaluate Librep expressions. Suppose for instance you want to stop the saw started by the `.sawfishrc` on the right. The command to do that is `3d-destroy`. From the command line type `sawfish-client` and the following prompt appears: `sawfish%` Just type in `(3d-destroy)` and the window with the rotating saw will disappear. Of course, this is just an example: `sawfish-client` can be used for more sophisticated purposes. The command `sawfish-client -?` will display the fundamental options of `sawfish-client`.

GETTING HELP

Unfortunately, there aren't any easy-to-follow tutorials included with *Sawfish*. The best way to browse the documentation that comes with it is to use the **info** utility. To do this, call up a terminal and type: **info sawfish** (**man info** for a detailed description of the **info** command). This will open an *Emacs* window that can be navigated with the arrows and the Enter key.

The documentation here is intended as a reference more than as a tutorial and can be difficult to read for beginners, although this article should introduce some of the basic concepts.

You'll find some more documentation in the directory `/usr/share/doc/sawfish-0.30.3` (where 0.30.3 is the version number of Sawfish you have installed). Particularly useful here is the FAQ section, that should answer almost all the basic questions and give a few hints on where to find more information. The best way to get information for Librep is to type: **info librep**. This reference manual contains a decent introduction to the programming language.

The other option is to hit the 'net and make use of others' experiences. Check out www.sawfish.org for more.

→ add the line **(require 'sawmill-defaults)** which will force *Sawfish* to load sawmill-defaults.jl even if there is a .sawfishrc file. Simple Librep code can be included directly into .sawfishrc, but if you plan to heavily customise *Sawfish* by adding a lot of extensions, it is a good idea to put the extra functions in ~/.sawfish/lisp and to load them by adding a **(require 'xxx)** statement to your .sawfishrc file.

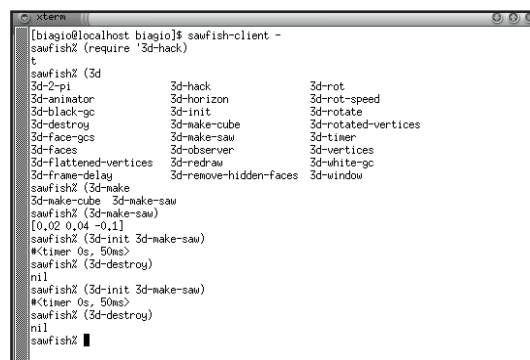
Here is an example of an easy-to-understand `.sawfishrc`:

```
; Load the sawmill-defaults file
(require 'sawmill-defaults)
; A toy :)
(require '3d-hack)
; (3d-init 3d-make-cube)
(3d-init 3d-make-saw)
; Load my extensions
; The following code written by Peter Makhholm
<peter@makhholm.net>
; avoid duplicated names of windows
(require 'number-windows)
```

The lines starting with a semi colon are comments. It is a good idea to comment your configuration file, since it allows you to remember the meaning of each line. The first line interpreted by *Sawfish* loads the sawmill-defaults file. This action avoids loosing the default settings.

The line (**require '3d-hack'**) loads a library distributed with *Sawfish*. This library defines two functions that creates three-dimensional rotating geometrical figures on the desktop. The line (**3d-init 3d-make-saw**) creates a saw, while the commented line (**3d-init 3d-make-cube**), if called, creates a cube. These lines give an example on how to insert simple Librep code inside a .sawfishrc file.

Finally, the private library **number-windows.jl** (located in `~/sawfish/lisp`) is loaded. This library (downloaded from <http://www.sics.se/~lofgren/sawmill/>) avoids duplication of names in the title bars of similar applications. For instance, if I open two terminals, the first will be named 'Terminal' and the second 'Terminal<1>'. In this way, they can be referred in the Windows section of the menu with different names. We won't go any deeper into the source code of that function. It's just to reinforce the idea that complicated functions should be put in the private repository and called via a **require** statement in the `.sawfishrc`. This makes `.sawfishrc` easier to read.



Sawfish-client executing some Librep code described in the text. The appearance of the xterm is defined by the theme microGUI, which reproduces the look of the QNX operating system.

THE CONFIGURATION UTILITY

There is a third customisation file: **custom**, which resides in the directory `~/sawfish`. This file contains the customisation options generated by the configuration utility of *Sawfish*, *sawfish-ui*. Since sometimes the order matters (before using a function you have to define it), it is worth noting that **custom** is loaded after **sawmill-defaults**, but before **.sawfishrc**. As **custom** is automatically updated, it is a good idea not to edit it by hand; just use *sawfish-ui* to modify its content. The configuration utility can be invoked either from the *Sawfish* menu or typing in an xterm: **sawfish-ui &**. The interface provides options to customise almost every aspect of the window manager, and is divided into the following sections:

Appearance. Changing the general look and feel of *Sawfish*.

Bindings. Defining key bindings (i.e. bindings of Libreps commands to a given combination of keys).

Focus. Setting the properties of focused and unfocused windows.

Matched Windows. Defining the actions to be taken for a particular class of window (e.g. the xterm).

Minimising/Maximising. Setting the characteristics of minimised and maximised windows.

Miscellaneous. Changing minor categories like Tooltips.

Move/Resize. Defining the properties of windows when these actions are taken.

Placement. Setting options for the placement of windows.

Sound. System sound setup.

Workspaces. Configuring navigation through virtual desktops.

All these options are very easy to use. Here's an example: suppose that we want to open an xterm with a combination of keys. After opening the section 'Bindings', we choose the context 'Global' (to which the function that open the xterm belongs) and we click on 'Insert..'. We get the list of all the functions that belong to the context 'Global'. Then select from the list 'xterm'. The next step is to bind a key to the action. If we click on 'Grab', we will be asked to press a combination of keys on the keyboard. We then press **x** while holding **Alt** and click on OK. From this moment on, every time we repeat that combination of keys we open a new xterm.

The automated configuration utilities of *Sawfish* can also be accessed from the Control Centre in *GNOME*.

Themes

It is impossible to deny that appearance is often the characteristic that attracts users to a specific window manager. *Sawfish*, like a lot of others, allows the user to change most aspects of the appearance.

The appearance is determined by a theme, a generic term that seems to have a different meaning depending on which window manager you happen to be using. In the case of *Sawfish*, a theme is composed by the image of the titlebar, the images of the maximise, kill, iconification and option listing buttons and the fonts of the titles of the windows, but not the desktop background!

The package is usually distributed with a number of themes which should give you a good idea about how to set up your own. Changing the current theme is simple using *sawfish-ui*. After launching it, select 'Appearance' from the left menu. A window will be displayed that allows you to change, among other things, the default frame style. You can select the one you like by browsing all those you have installed. The other parts of the right pane allow you to change some default settings of the theme.

One of the neatest (but slightly pointless) features of *Sawfish* is that it is possible to select a different theme for every open window: just right-click on the title bar of a window and select a theme from the 'Frame style' menu. Only the selected window will show this theme.

Fortunately you're not stuck with the supplied themes, as new ones can be downloaded from <http://sawfish.themes.org> and then installed. We'll give a few examples of how to install new themes. However, be aware that the following example worked on our Mandrake 7.2 distro, but doesn't correspond with the theme how-to located in sawfish.themes.org. This default installation procedure scattered error messages all over the place!

Generally speaking, a theme can be installed for personal use, or as root, can be made available to all users. In the first case, the tar ball with the theme must be located in `~/sawfish/themes`. Suppose we downloaded the file `sawfishtheme.tar.gz`. According to the general instructions, we store it in `~/sawfish/themes` and after a restart of *Sawfish*, it should be available in the theme selection window. However, when we tried to select the new theme, it didn't work: we had to decompress it with the command:

```
tar xvzf sawfishtheme.tar.gz
```

after changing the current directory to `~/sawfish/themes`. Then, after another restart, it was available.

For public use, the tar ball has to be stored in `/usr/share/sawfish/themes`. Also in this case, there were problems with the `tar.gz` file, and we had to convert it in a `tar.bz2` file (the default and the alternative procedure were repeated on more than one tar ball, with similar results).

These problems aren't too trying; *Sawfish* is a relatively young application, and the automatic procedures can still have some problems. Moreover, since it allows powerful customisation, it could be that the version installed in

our system has some different features from the original one. However, in the case of installing a new theme, the decompression of the tar ball should work in most cases.

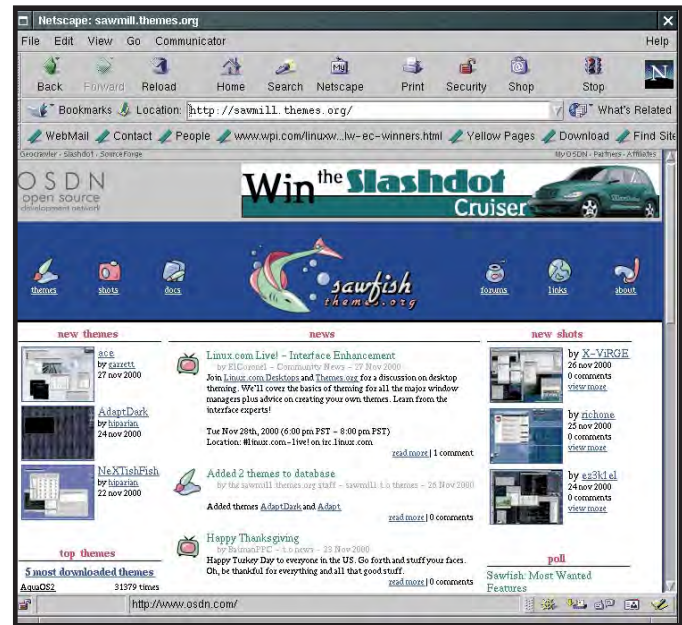
DO IT YOURSELF

If you are still unhappy with the various themes you've found on the Internet, you can always create your own. For this purpose, *Sawfish* makes available the utility *sawfish-themer*, distributed with the latest release. The subject of theme creation is so broad that we don't have time to cover it in this short tutorial, but some useful information can be found in the FAQ on the project's homepage.

But what about the background? Is there any way to set it direct from *Sawfish*? Of course there is, thanks to Librep. Just add the following line to your `.sawfishrc`: **(run-shell-command "xv -root /home/biagio/IMAGES/LinuxSleep.jpg -quit")** changing `/home/biagio/IMAGES/LinuxSleep.jpg` with the complete path to the image you want to see. Incidentally, the above example shows how to interact with the shell in a Librep program. Of course, if you do not have `xv` (which can be downloaded from <http://www.trilon.com/xv/downloads.html>) you can replace the command line with any other that allows to set a graphic file as background.

A LAST WORD

Sawfish is a relatively new project that has serious potential. Its greatest characteristic is the extensibility brought about by its modular structure. In this article, we have described the first steps for customisation, mainly by looking at the available tools. For more direct, drastic alteration, you'll need at least a basic understanding of the Librep programming language. We have introduced some features of Librep from a practical point of view, and by looking to the examples here you should have no trouble personalising *Sawfish* beyond the possibility offered by the automatic tools. For those who want to really get into the guts of the application, you'll have to wade through the documentation. **LXF**



New themes for Sawfish can be downloaded from <http://sawfish.themes.org>.

THE APROPOS COMMAND

The **apropos** command is the easiest way to get documentation for Librep options. A way to call it is to bind a key to **apropos** in the Global Context (Binding option of *Sawfish-ui*). Then, when the corresponding key combination is hit, a window with a request to insert the name of the function will appear. Enter the name of a function and the documentation will be displayed.

PERLS

OF WISDOM

Regular expressions and Packages



FIND OUT MORE

Note that this is a really abbreviated intro to regular expressions. For the full story, you need to read Chapter 3 of *Programming Perl, 3rd Edition* by Larry Wall, Tom Christiansen and Jon Orwant, ISBN: 0-596-00027-8, or *Mastering Regular Expressions* by Jeffrey E. F. Friedl, ISBN 1-56592-257-3.

Charlie Stross continues his exploration of Perl with a look at how the language manages text and its mechanisms for enforcing namespaces.

One of the most interesting features of Perl – which we don't find built into any traditional languages like C or Lisp – are regular expressions. We use regular expressions for scanning text strings and carrying out search/replace operations on them. It's possible to use Perl for text mangling without regexps (as they're known) – Perl also provides the BASIC **substr()** and **length()** functions – but regular expressions are at the heart of much of the language's power.

If you've used a word processor's search and replace facility, you're used to entering a word and having the word processor search for the first occurrence of that word in the text. You may also have used a word processor that understands patterns: a series of characters that represent items that can be found in the text. Regular expressions in Perl are basically a sophisticated method of specifying patterns in text.

KEEP IT SIMPLE

About the simplest kind of regular expression is one that looks like this:

```
/red/
```

This isn't just the word 'red' enclosed between slashes: it's a pattern that matches the character 'r' followed by the character 'e' followed by the character 'd'. In a regular

expression, any character (with some exceptions we'll see in a moment) is a pattern. Simple letters or integers usually match themselves, but most punctuation characters have special meanings. We use this regular expression like this:

```
$text = "the quick brown fox jumped over the red
hen";
if ($text =~ /red/) {
    print "caught red handed!\n";
}
```

The operator **=~** 'binds' a regular expression to a variable (in this case, to **\$text**) and returns true if a match exists between the pattern and the contents of the target variable. It also sets some special variables (see box, top right).

So, if we want to delete everything up to and including the matched text, we can do something like this:

```
$text = "the quick brown fox jumped over the red
hen";
if ($text =~ /red/) {
    $text = $';
    print "$text\n";
}
```

which will print 'hen'.

You can also extend their use with a series of metacharacters, see box, right. So, to match primary colours we can say something like:

```
$text =~ /(red)(blue)(green)/i
```

Note the trailing **i** after the last slash. This is known as an expression modifier; modifiers are used to change the way the Perl regexp engine interprets an expression. This one means that the regexp is to be evaluated as case-insensitive, i.e. as either upper or lowercase, or mixed case. It saves us from having to write something like:

```
$text =~ /([Rr][Ee][Dd])([Bb][Ll][Uu][Ee])([Gg][Rr][Ee][Ee][Nn])/
```

There are a number of other modifiers, however they're specific to the type of pattern we're using.

REPLACING STRINGS

There are, broadly speaking, three ways you can use a regular expression. You can deploy one (as above) to match (search for) a pattern to a string. You can use one to match one pattern and replace it with a different one. And you can use it to match (and replace) a set of characters.

The form we've been using (`$variable =~ /expression/`) is a simplified version of the **m//** (match) operator. To match and replace we use the **s//** (substitution) operator, and to operate on a set of characters we use the **tr///** (translation) operator. Simple replacement looks like this:

```
$text =~ s/red/blue/;
```

This replaces every occurrence of the string 'red' with the string 'blue'. Each matched sub-expression is in turn assigned to special variables called `$1`, `$2`, and so on:

```
$text =~ s/red/blue-$1/;  
print $text;
```

will print 'blue-red' if successful.

Note that it might not be what you want to achieve – instances of 'redundancy' will become 'blueundancy'. In fact, if we want to replace the word 'red' with the word 'blue', we have a couple of implicit problems to solve; how to identify word boundaries, and how to apply a change to a subexpression. Perl provides some special metasympols that help here. We've seen **^** used to match the beginning of a string: For example, if we do:

```
$text =~ /^red/
```

it only matches if the string stashed in `$text` begins with 'red'; **red\$** does the same for a string ending in 'red'. There are a number of metasympols available for string matching, (see box, above right) for a small selection.

Our problem of replacing 'red' only where it's a word can now be solved. First, we bracket 'red' with metasympols that

VARIABLES

VARIABLE	CONTENTS
<code>\$</code>	All the text in the target string, to the left of the match.
<code>\$&</code>	The text that matches the regular expression.
<code>\$'</code>	All the text to the right of the matched section in the target.

METACHARACTERS

Metacharacters are special characters that don't represent themselves (in a regexp); they form a simple language.

CHARACTER	MEANING
<code>\</code>	The following character is not a metacharacter (e.g. <code>\\</code> is treated as a literal <code>"\"</code>).
<code> </code>	Logical-OR; <code>/a b/</code> means 'match either a OR b at this position'.
<code>()</code>	Group some characters together to form a sub-expression: <code>(red)(blue)</code> means 'match either 'red' OR 'blue', not 'match 'red', then '(' or ')', then 'blue'.
<code>[]</code>	Character set: <code>[A-M]</code> means 'match any one character in the series [A, B, C ... M]'.
<code>^</code>	Outside a character set, this means 'match the beginning of a string'.
<code>^</code>	As first element in a character set, means 'match anything that is not a member of this set': <code>[^A-M]</code> means 'match any character that is not in [A-M]'.
<code>\$</code>	Matches the end of a string.
<code>*</code>	Match zero or more of the preceding expression.
<code>+</code>	Match one or more of the preceding expression.
<code>?</code>	Match exactly zero or one of the preceding character (equivalent to <code>{0,1}</code>).
<code>? (other)</code>	If it occurs after a quantifier (like <code>"</code>), match the shortest possible string that the preceding expression binds to. Match any single character once.
<code>{a,b}</code>	Match from a to b occurrences of the preceding expression; if 'a' is missing, match up to 'b' occurrences, and if 'b' is missing, match at least 'a' occurrences. e.g., <code>/(red){2,3}/</code> means 'match two to three occurrences of the expression (red) in a row'.

METASYMBOLS

METASYMBOL	MEANING
<code>\d</code>	Match any digit character.
<code>\D</code>	Match any nondigit character.
<code>\w</code>	Match any word character ([A-Za-z_]).
<code>\W</code>	Match any nonword character.
<code>\s</code>	Match any whitespace character (space, tab, return, etc).
<code>\S</code>	Match any non-whitespace character.
<code>\3</code>	Match the third previously-matched subexpression (and <code>\2</code> , and so on).

There are other metasympols; they're listed in table 5-7 of Chapter 3 of *Programming Perl*.

match only nonword characters, so that a word like 'redundancy' won't match – the 'u' is a word character. This gives us an expression like:

```
$text =~ /\W+red\W+/
```

Which only matches the word 'red', not words containing the substring 'red'.

Then we do the substitution like this:

```
$text =~ s/^(.*\W*)red(\W+.*)*$/\1blue$2/
```

➔ Let's break this down. The search pattern (between the first two slashes) has a subexpression (grouped with brackets) on each side of the word 'red', which we're trying to replace.

The first subexpression is anchored to the start of the string (with the caret). It looks for the start of the string, followed by a sequence of zero or more of any character, followed by zero or more non-word characters. (That's because we might have the word 'red' right at the start of the string.) We then have the word red – but this won't match if red isn't preceded by a non-word character, as in 'inferred'.

The second subexpression states that 'red' must be followed by the end of the string – optionally preceded by a group containing one or more nonword characters, then some more arbitrary characters.

On the right hand side of the replace expression, we see **\$1blue\$2**. What this means is, the match (which starts with ^ and ends with \$, and therefore amounts to the entire string, if – and only if – it matches) is to be replaced with the first bracketed expression, the string 'blue', then the second bracketed expression. And both of these expressions include the nonword stuff that precedes (or follows) our target word.

Incidentally, the **s///** (replace) operator takes a modifier that makes it a lot easier to read: if there's a trailing **/x**, we can include whitespace and comments in our regexps:

```
$text =~ s/
    ^("W") # group match leading
    chars and whitespace, if any: $1
    red    # match literal text
    "red"
    (W+.)$ # group match trailing
    whitespace and text, if any: $2
/
    $1blue$2
/x ;
```

If you find the slashes confusing, you can replace them with braces like this:

```
$text =~ s{
    ^(          # open group $1, anchored to
    start of string
    ."W"       # match chars and whitespace, if any
```

```
)
red          # match literal text "red"
(           # open group $2, immediately
after "red"
\W+.*      # match whitespace and optional chars
)*         # group two is optional (zero or
more occurrences)
$          # anchor to end of string
}
$1blue$2   # what we replace this with
}x;
```

REPEATABILITY!

There is one fundamental problem with the regular expression we built up in the last section: it doesn't work properly if there are two or more occurrences of 'red' in a line. For example, 'the red house has a red door' will be replaced by 'the red house has a blue door'.

The reason it doesn't work is that regular expressions are greedy – they will match as many characters in the string as possible. The initial grouped expression, \$1, can match either 'the' or 'the red house has a'. Being greedy, it will take the longer string; thus we'll end up only replacing the second 'red'.

The way we get around this is to remember that a regular expression match using the **s///** operator is like a function; it returns either undef (zero) if no matches are found, or the number of matches it made. So we can put the whole thing in a little loop:

```
1 while $text =~ s/^("W")red(W+.)$/$1blue$2/ ;
```

The statement **1** doesn't do anything, but it acts as a body for a **while** loop, the condition of which is the search/replace expression. While the search/replace returns a non-zero number, we have stuff to do, so Perl evaluates the non-functional statement '1' and executes the loop condition again; when there are no more replacements, execution leaves the loop.


Of course, this is the hard way to replace words in Perl: as the language motto says, 'there's more than one way to do it'. In the table of metasymbols above, we met symbols that represent word characters or nonword characters. But there's a different set of symbols to represent positions (see box, left).

So, with these appositional symbols we can re-write our substitution as:

```
1 while $text =~ s/^(.)bred\b(.*)$/$1blue$2/ ;
```

TRANSLATING SETS

In addition to searching and replacing strings, Perl's regular expression engine can be used to translate characters. The operator we use for this is **tr///** (transliterate), and instead of working on regular expressions, it works on characters: **tr/a/b/** replaces all occurrences of the character 'a' with the character 'b', for example. More usefully, **tr///** works on sets of characters: **tr/[a-g]/[h-n]/** replaces a's with h's, b's with i's, c's with j's, and so on (but doesn't do anything to characters later than g).



APPOSITIONAL METASYMBOLS

Symbol	Meaning
\b	Word boundary (the position between a \w and a \W, in any order).
\B	Matches any position that is not a word boundary.
\A	Beginning of string (even if multi-line matching is switched on).
\z	Matches at end of string.
\Z	Matches right before new line at end of string (or EOS if no newline).

tr/// MODIFIERS

Modifier	Meaning
/c	Complement search list
/d	Delete characters that don't have a defined replacement
/s	Squash duplicate replacements

Here's a simple one:

```
$text =~ tr/[a-z]/[A-Z];
```

This translates all lowercase letters in `$text` into uppercase.
And here's another:

```
$text =~ tr/A-Za-z/N-Za-Mn-za-m/;
```

This is the classic Caesar substitution cypher – also known as ROT-13 – where every character is replaced by one shifted halfway through the alphabet. (Note that the square brackets previously used to identify a set are redundant in a transliterate expression.)

A useful mechanism to note is that you can represent any character in the 8-bit ASCII codeset as a backslash followed by an octal number (three digits, padded with leading zeroes). For example: `tr/003/C/` will translate all occurrences of Control-C into a capital C. It's also worth knowing that as of Perl 5.6, Perl understands POSIX character classes – a set of specially defined classes with names, such as `[:digit:]` (equivalent to `[d]`), or `[:space:]` (equivalent to `[s]`). Moreover, Perl has the basics of Unicode support – but that's a long, hairy, topic that is best left to Chapter 3 of *Programming Perl*.

There are some handy modifiers to `tr///` (see `tr///` Modifiers box, left). For example:

```
$text =~ tr/a-z/a-m/d ;
```

will delete all letters from n-z inclusive from `$text`.

PERL ONE LINERS

Perl isn't just good for big programs; you can also use it as a power-tool on the command line:

```
perl -e 'print "Hello, world\n";'
```

We can type a short Perl script after the `-e` flag and Perl will execute it; make sure that you single-quote it, and put a backslash in front of any double-quotes.

The `-n` flag is the most useful. It makes Perl execute the code you specified after the `-e` flag in a loop, iterating over each line read from standard input. If we say:

```
perl -ne '$i++; print "$i: $_";'
```

what is actually executed is equivalent to:

```
#!/usr/bin/perl
while (<STDIN>) {
    $i++;
    print "$i: $_";
}
```

(print everything read from standard input, prefixed with a line number). A piece of Perl magic comes in handy here: if we say something like: `perl myprog.pl foo bar quux`, Perl will run `myprog.pl`, and attempt to successively read the files `foo`,

MODULAR PROGRAMS

When we use a variable like `$text` or define a subroutine called `&mysub()`, Perl stores it in a chunk of memory called a namespace. By default, everything lives in a namespace called `main`, and if you refer to `'$text'` without specifying the namespace, Perl assumes you're talking about `$main::text`. We can define variables in other namespaces; they spring into existence as soon as we mention them: for example `$myprog::text = 1`.

`$myprog::text` and `$main::text` are two completely different variables. They don't know about each other or share contents. Likewise, `&main::mysub()` may be very different from `&myprog::mysub()`.

Writing down the namespace every time we want to define a variable or subroutine in a foreign one can get tedious, so there's a simple keyword that tells Perl we're in a new namespace: `package()`. If your Perl program contains a line like this:

```
package MyPrivate;
```

Every variable and subroutine defined after it will be prefixed with `MyPrivate::` until you exit the package. (You exit a package by exiting the enclosing scope – end of file, end of `eval()` block, or end of a bracketed chunk of code.)

Note that it's important to end a package with some expression that evaluates to 'true' (we'll see why in a moment). Thus, it's common to end a package with `1;` (one, being non-zero, is therefore true all the time).

The useful thing about packages is that they allow us to keep a bunch of code

and data somewhere private – this is important when writing a large, modular program. However, the annoying thing about packages is keeping track of them, and of the namespaces they use. A common, conventional solution is this: each package goes into a separate file, which begins with the line `package <somepackagename>` and ends with `1;`. The main program then has a series of lines at the top:

```
require "package1";
require "package2";
# and so on
```

The keyword `require` tells Perl to search the directories in its module search path (a list of directories in the internal variable `@INC`) for a named package. You can either use the name you defined in the package file, or specify a relative pathname to the file. When Perl finds the package file, it loads, compiles, and executes it – which is why we want it to return '1' (otherwise we'll end up with a runtime error). Usually we just put subroutines and variables in packages, so nothing obvious happens – but the extra namespace and code is sitting around waiting for our main program to use it.

There's a related keyword, `use`, which is similar to `require` except that it is executed while your program is being compiled, instead of at run time; this means that errors will be detected sooner, and it allows you to set up your variables and package code before running the actual program.

bar, and quux on standard input. (But Perl still understands the usual shell input/output redirection conventions.) We can use this for editing the contents of files:

```
perl -ne 's/red/blue/g; print; ' fred.txt
```

(Transform the string "red" to "blue" throughout `fred.txt`; print the result to standard output.)

This is great for writing short filter scripts, but what if we want to edit a bunch of files and save the original copies? The `-i` flag is our friend; specify `-i.bak` on the command line, feed your scriptlet a bunch of files on standard input – and instead of sending the output to the standard output, Perl will add the suffix `.bak` onto each input file and save the output in a file with the original name. So:

```
perl -i.bak -n -e '$i++; print "$i: $_";' article.txt
```

saves a copy of the file `article.txt` as `article.txt.bak`, and writes a new copy with leading line numbers. **LXF**

NEXT MONTH

In next month's tutorial we'll see how packages can be used to create objects – bits of data that have associated private code – and where `use` fits into the scheme of things.

SPEAKING JAVA

Part 6: Introduction to threads



Java can think about more than one thing at a time, as Richard Drummond explains.

We expect programs to be able to do many tasks simultaneously. Consider your web browser, for example. When you open a web page, it fetches a page from the 'net, typically laying out the page as it does so and simultaneously fetching any images on the page as well. While all that is happening, you still expect the browser to respond to your input, whether to cancel loading or browse to another page. The browser is able to do all this because it is a multi-threaded application. Within the one program, there are multiple threads of execution running at the same time. Traditional programming languages have no built-in support for threading, so writing a multi-threaded program would usually mean implementing it in some platform-dependent way. Threads are integral part of Java, however, and their use is supported by the standard class library itself.

WHAT IS A THREAD?

A thread, then, is a kind of light weight process – with less overhead than the usual processes of a multi-tasking operating system – that exists within the context of a single program. In any Java program there are typically many threads running at a given time, even if you don't explicitly create your own threads. The garbage collector in Java, for example, which runs in the background and frees the memory associated with unused objects, is a thread.

Like processes, threads have a sequence of flow. They have a beginning, a middle and an end. Your main program is run as a thread, beginning at the **main()** method of your top-

level class and terminating when it reaches the end of that method or when an untrapped exception occurs. A Java program will terminate only when all of its threads have been terminated.

A Java thread is handled through the **java.lang.Thread** class abstraction. Every thread has an instantiation of this class to act as a handle on it. To create a new thread, then, we simply create a new **Thread** object.

CREATING THREADS

There are two principal ways to structure the creation of threads. The first, and perhaps the simplest, is to directly subclass the **Thread** class, and this is what we'll do in our examples here; the other method is one we'll look at next time, and that is to create a class which implements the **Runnable** interface and to pass that class to a **Thread** object's constructor when you create it.

Now, the **Thread** class has a method called **run** which is where all the action takes place. That is, it contains the sequence of instructions that will be executed within instances of that **Thread**. In the **Thread** class itself, this method is empty, and starting an empty thread will do nothing. In our sub-class, however, we override it to create the body of the thread. Have a look at this example.

```
public class SimpleThread extends Thread
{
    public void run()
```



```
File Sessions Options Help
bash-2.04$ java ThreadDemo1
Thread-0 is running
main is done.
Thread-1 is running
Thread-1 is running
Thread-1 is running
Thread-1 is running
Thread-0 is running
Thread-0 is running
Thread-1 is running
Thread-1 is running
Thread-0 is running
Thread-1 is running
Thread-1 is running
Thread-0 is running
Thread-1 is running
Thread-1 is running
Thread-0 is running
Thread-0 is running
Thread-0 is running
```

Fig1: This output was generated by running two instances of the class *SimpleThread1*.

```
{
for( int i=0; i<10; i++ )
{
System.out.println( getName() + " is running" );
try
{
sleep( (long) (Math.random()*2000) );
} catch (InterruptedException e) {}
}
}
```

Threads created from this class will loop ten times printing out the name of the thread that it is being executed as (the `getName()` method returns the name of a thread) and then terminate. The `sleep()` method used above puts a thread to sleep for a specified amount of time in milliseconds. Here, the thread sleeps for a random period of up to two seconds. We'll explain why we've done this in a moment. First we need a class to act as a driver and use some threads created from our example class.

```
public class ThreadDemo1
{
public static void main( String[] args )
{
new SimpleThread().start();
new SimpleThread().start();
}
```

```
System.out.println( Thread.currentThread().getName()
+ "is done.");
}
}
```

Threads are created like any other object with the **new** keyword. Simply creating a **Thread** object does not cause the thread to begin executing, however. We need to invoke its **start()** method for this to happen. So that's what we do here. We create two threads and start them off. If you run *ThreadDemo1*, you should see an output something like that shown in *fig1*.

There are a couple of things to point out here. Firstly, we haven't explicitly named the threads we have used here – you can do this via the **Thread** constructor or with the **setName()** method – so we get default namings for our threads, which is 'main' for the main thread, 'Thread-1' for the first thread we create, 'Thread-2' for the second, etc. Secondly, the **getCurrent()** call is a static method in the **Thread** class and returns a reference to the currently executing thread. And thirdly, as you can see from the output, our main thread is finished long before any of its two child threads have terminated, but the program carries on running, anyway. This might not have been the behaviour you expected, but is perfectly normal in Java. It is quite easy to change this, however. You can make one thread wait for another one to terminate with the **join()** method. Invoking **join()** with no parameters on a thread makes the current thread sleep until that thread has finished. For example:

```
public class ThreadDemo2
{
public static void main( String[] args )
{
Thread t1 = new SimpleThread().start();
Thread t2 = new SimpleThread().start();

t2.join();

System.out.println( Thread.currentThread().getName()
+ "is done.");
}
}
```

SCHEDULING

On a single processor system, obviously, the CPU can only execute one instruction at a time. Then how does it appear that more than one thread can be running simultaneously? Well, the executions of threads are interleaved – the processor jumps from one thread to the next. For example, the processor might execute the first instruction in thread 1, then the first instruction in thread 2, then the second and third instructions in thread 1, the first instruction in thread 3, and so on. Every time the processor switches from one thread to another, any state required to resume processing that thread later is saved. The end result is what we want: that the sequence of instructions in each thread are processed in order. We just don't know when an instruction in one

Stopping Threads

As a counterpart to the **start()** method, there is also a **stop()** method which forces the termination of a thread's execution. Its use has been deprecated, however, because it can be unsafe. It is much better for a thread to arrange its own demise and simply fall out the end of its method.

Priority

A thread's priority is set via the **setPriority()** method and read via the **getPriority()** method. The higher the priority of your thread, the more chance it has of running.

WIN



COMPAQ

iPAQ

each week*

To enter, visit
www.greymatter.com

www.greymatter.com

The serious choice for software
developers - delivering a full product
range at competitive prices with
strong technical support

GREY MATTER

supporting software developers

Tel: +44(0)1364 655 108
Email: maildesk@greymatter.com*For full terms and conditions, visit
www.greymatter.com/iPAQ

→ thread we'll be processed relative to an instruction in some other thread.

The job of interleaving the execution of the various threads running is done by the scheduler. It decides which thread is to get a chance to run next. This is based on the state of a thread and its priority.

When the **start()** method is invoked on a thread, it goes into the state 'Runnable'. This means that it is eligible for execution. It will stay in this state until it sleeps or until it terminates. A thread is put to sleep when it is waiting for some event to occur. The thread will be notified of the event, probably by a signal from some other thread, and become Runnable once more. An example of this is the **sleep()** method. This puts the thread to sleep waiting for a wake-up call from a timer.

In our SimpleThread class above, the thread gets put to sleep for a random amount of time for every iteration of the loop. This is saying: give some other thread a chance to run and wake me up in a bit. We have done this so that in the output of the program, the two threads are mixed up randomly. But what would have happen if we didn't do this?

Modify SimpleThread as follows and try running the example again (see *fig2*).

```
public class SimpleThread3 extends Thread
{
    public void run()
    {
        for( int i=0; i<1000; i++ )
        {
            System.out.println( getName() + " is running"
        );
        }
    }
}
```

What happens now? Well, it largely depends on what JVM you are using and what operating system it is running on. In Sun's JDK1.3, Java threads are implemented as native threads on the host operating system – on Linux this means light-weight processes – whereas earlier JVMs used so-called green threads.

Java only guarantees that a thread of the highest priority will be running. If that highest priority thread does not go to sleep at any point, it is possible that no other thread will get a chance to run. Some operating systems will implement some sort of time-

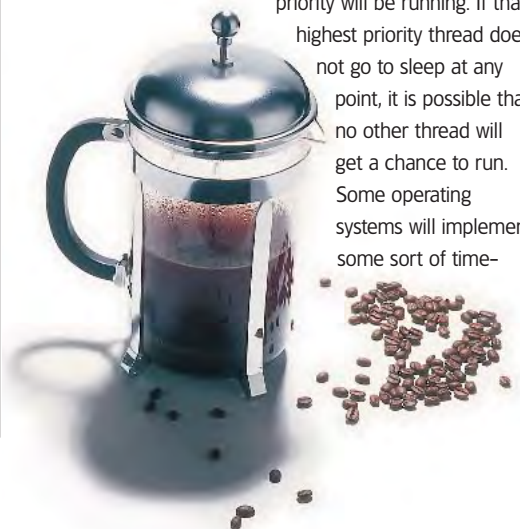



Figure2: The SimpleThread3 class implements a greedy thread which won't give up the processor voluntarily.

slicing, however. That is, if there are two or more threads of the highest priority, then one of them is executed for a fixed period, and if, after that period, has not been put to **sleep()** then the scheduler will pre-empt that thread and give another of the highest priority threads the opportunity to run.

The upshot of all this, is that when writing Java programs, you should not rely on the scheduling behaviour of one particular JVM. You can only guarantee that a thread will be pre-empted by one of higher priority; you cannot say that it will be pre-empted by one of the same priority or not.

As you can probably begin to tell, writing multi-threaded programs opens up a whole new world of problems to think about while designing programs. Next time, we'll look at the tools Java provides to let threads synchronise with each other, and we'll see what potential pitfalls that may cause. 



ANSWERS

If you are really stuck, why not write in? Our resident gurus will answer even your most complicated problems!

Once again we're here to answer your technical queries and conundrums. Linux can be tricky at times, and everyone needs a helping hand every now and then.

As I'm sure you've already discovered, Linux users are a friendly bunch, only too keen to help you with your problems. That extends to the magazine too, where we'll have a go at answering any query you send us, no matter how simple, now matter how difficult.

Please feel free to email or write in with your own particular queries. The most interesting ones will be printed here every issue. We won't be able to reply personally, so please don't be offended if you don't hear from us. If you want to stand a better chance of having your question featured, read the submissions advice boxout on page 91!

No question is too tricky for us to track down the answer, so do your worst – we've got a team of experts from around the Linux world to take on the challenge!

Nick Veitch
Editor

Double drive dilemma

Q I recently purchased a copy of *Linux Format* with the Corel Linux CD on the front. I've been wanting to try Linux and think I will install a 4.3Gb extra hard drive, put all the Linux applications on there and run a dual boot system. I have not attempted the install yet, as I want to be sure that I can run both Windows and Linux on different hard drives although not at the same time. I want to make sure I don't mess up all the software I presently have and use. Are there suggestions you can give me?

A Windows is rather picky and needs to be installed on the first partition of the primary master drive in the system. Putting aside that restriction, you can distribute your Linux partitions as you choose. You might want to put your second hard disk on the second IDE controller and put your swap partition on the first hard disk, so swap and the normal filesystem don't have to fight for IDE bandwidth all the time.

Other than that, you'll want to install LILO on `/dev/hda`, rather than `/dev/hdb` or `hdc`,

so that booting the machine will give you the LILO menu where you can select the system to boot. Most distributions will allow you choose where LILO is installed, so you won't need to edit any configuration files later on.

Different class

Q I think I have installed the wrong installation class of Mandrake 7.2 (*Linux Format* issue 09 coverdisc). I installed the 'recommended' class, but I now realise this installation class does not include all the programming languages and other developer stuff. In order to install a 'workstation' installation class, do I need to remove my current installation, or can I simply update it, by

installing the 'workstation' installation class over the top?


What is the difference between the covermount versions of Linux and the commercially available boxed editions?

And finally, do you know if there is a good, stable Atari ST emulator available for Linux?

A You can use *LinuxConf* to install additional packages for the items you missed out, or you can install them by hand using RPM. You might also want to try to 'upgrade' the distribution using the Mandrake installer, as long as it doesn't notice that your already running the current release and stop you.

Boxed sets usually contain CDs containing the actual distribution, the source code and some

Deb? RPM? Can you turn one into the other?



variety of printed documentation. As well as that, support is available from the distribution creators, such as RedHat or SuSE, for a fixed time period after installation should you become stuck. Boxed sets are good for people who've never touched Linux as they have something to follow when they install, but as the distribution can be, er, distributed freely, it's useful for people who can obtain help elsewhere to install Linux off a free, or cheap, CD rather than shelling out on a boxed set which they'll hardly use.

On the Atari front, a quick search with Google brings up a few hits, the most promising of which seems to be <http://www.complang.tuwien.ac.at/nino/stonx.html>, which is the home of STonX. Another good site is <http://www.atari.org/links/ST/Emulation/Emulators/> which has a number of interesting leads. We've not tried any of these out, so we can't promise anything... The alternative is to get your old ST out of the loft and dust it down.

Debs and RPMs

Q I am using Linux-Mandrake 7.2 and I am very happy with it. However, I would like to be able to install Debian packages on my system e.g. smartdrive. Is it possible to install dselect alongside rpm-drake and kpackage. I have tried copying dselect from the Corel Linux disk but my system won't have anything to do with it. I would appreciate any advice you may have on this.

A It is very difficult to migrate a RPM based system to apt, so you may just want to reinstall with Debian 2.2 and start from a clean system. I imagine running dual package systems is doable, given the time and understanding of how apt works, but the time and effort required would outweigh the resulting improvements. Or you could add a Debian-based distribution to your system (see last month's MultiLinux tutorial for more).

Also, dselect has been replaced with apt in newer distributions of Debian, but as Corel is basically

Debian 2.0 with extra bits, they've not yet caught up.

Sound advice

Q I'm a newbie to Linux and recently installed Mandrake 7.2 to a partition of my PC's hard disk after discovering your mag at a newsagents.

I was amazed at how easily it installed. I was also surprised at the various window managers and packages such as KOffice that came on the disc. My problems may sound silly but I welcome your advice. They relate to the following:

I can't get the ESS 1868 Soundcard to work even though Mandrake recognises it, Linuconf states that the device is busy.

Also when I installed Corel Draw 9 it would not run when I click on the icon. I get the same problem with Opera 4 and OpenOffice when I unzip or install them. I installed them to their own folders which I created. (I logged on as a user as I had read it was dangerous to repeatedly log on as root).

A For the ESS card, which is supported by the Linux kernel, you will need to ensure that the isapnp system knows how where the card is, and that the ess1868.o module is passed the correct IRQs and IO addresses. For more details, as well as a good description as how to use isapnp read linux/Documentation/sound/ESS1868 within the kernel source tree.

For the programs which don't run correctly, you'll have to check what the icons are trying to run, then ensure you can run them from a terminal using the same command. As you didn't install them as root in a standard location, the binaries will not be in a standard search path for programs, so just running **opera** will fail. Passing the full path to the binaries should allow them to work correctly.

High anxiety

Q I'm still fairly new to Linux and taking the straightforward route during installation have managed to

successfully install, at different times, SuSE 6.4, SuSE 7.0 and a version of Slackware on a friend's computer. Partly due to your good review of Debian 2.2, I decided to give this distro a shot. The installation as you reviewed is problem free until I hit 'anXious' when it cannot determine my graphics card. However with SAX on SuSE configuring 'X' is

a breeze, so I know my card will work under Linux.

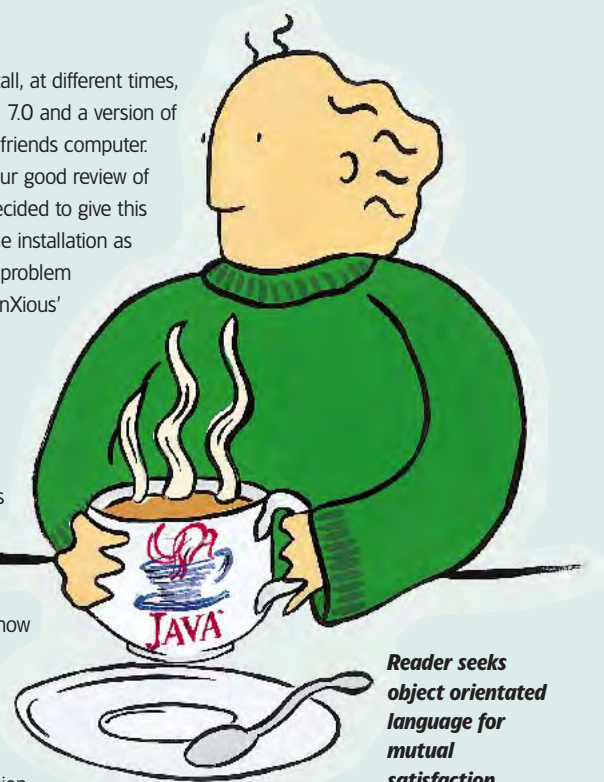
Once anXious on Debian has said 'I don't like your card', I feel the installation process became derailed because I'm not given the opportunity to manually tell Debian that my card is an ELSA ERAZOR III LT. How then do I tackle the job of installing X on Debian?

A Once Debian is installed you can reconfigure your X setup using 'XF86Config' as with any other distribution. You will want to ensure that you are running the latest version of XFree86, although you may be better off sticking with 3.3.6 rather than 4.0.2 due to various incompatibility issues with some popular cards. Once you've configured your card using XF86Config, you shouldn't have any other problems with it under Debian.

Java: there or not

Q I have an enquiry regarding Mandrake 7.2. I just bought a Linux Format magazine with the Mandrake 7.2 CD-ROM on the cover. I don't know if a Java compiler (jdk) was already installed on this distribution or do I need to download the Java software from Internet? Thanks.

A You can check to see if jdk is installed with the command



Reader seeks object orientated language for mutual satisfaction.

rpm -q jdk. If it is not there, check the Mandrake CD to see if it's on there before downloading the rather huge RPM from the Internet.

Small console-ation

Q I am very new at this and have only started using Linux today. I installed Linux Mandrake 7.2 from your issue 9 cover CD. Everything went well until Linux stopped loading up KDE (the thing where you choose which user name to use and whether to use GNOME, KDE etc.). I was wondering how to get back to this screen or to get KDE to load up. It just keeps loading up to the DOS 'type' screen, which isn't very helpful for new users.

A If you login from the text mode using the root user and the password you supplied when it was installed and open '/etc/inittab' in a text editor such as emacs or vi, you can change the line reading:

```
:id:3:default:
```

to

```
:id:5:default:
```

and save it. If you then do **init 5**, it should then start up with a →

→ graphical login, which will be the default for the next time you reboot the machine.

Norwegian Blue

Q I have two small questions for you about recent *Linux Format* cover CDs:

1. I've tried *OpenOffice* (from the December cover CD) but had to go back to *Staroffice 5.1* because I could not get the printer installed – that is, I could not print anything or correct the print format without having a printer installed, and there was no way I could find to install one. Any hints?

2. I tried *Mandrake 7.2* (from the Christmas CD) but went back to 7.1 because of a rather strange 'feature': when I pressed the @ key the mail program started... and as this key was used for this purpose I could not enter any email addresses – the key would not write an @. I use a Norwegian keyboard. Tried to change the keyboard to US, but this did not help, and anyway would have rendered the machine useless to me. Hints on this one? Anyway I didn't find *KDE2* that indispensable, and it slowed down my machine (AMD 350/64) quite a bit. It wasn't worth it.

By the way I've found a browser you seem to have missed in the browser-round-up:

it's called *Browse-X*. Small, fast, stable, terrific actually. Good to have when you're doing several things at once. Not very good for downloading big files, though it's much more of a finished product than *Arachne* for Linux. See it at <http://www.browsex.com>.

A The simple, if unsatisfactory answer to this is that you can't print directly from *OpenOffice*. It is very much a work in progress, and the initial release which went on the CD, was basically the source to the old *Star Office* project tweaked and with the proprietary elements removed. One of the non-free technologies licensed for use in *Star Office* is the printing part of the software, so this doesn't exist.

The *OpenOffice* project is an ongoing thing. You can check out the progress on the printing side of things at printing.openoffice.org

2. Your second problem does sound very weird, and doesn't sound like a keyboard mapping problem, so changing the keymap wouldn't solve it. It seems that something has set up a hotkey for that key to launch your mail software. This isn't part of the defaults for *Mandrake* or *KDE* – did

you install any other software after *Mandrake 7.2*? If it is something to do with the *KDE* settings, you can check out the keybindings in the Control Centre's "lookNFeel" section.

Thanks for the browser tip, it certainly seems that there is far more activity going on in the open source browser community than even we had imagined!

Code conundrum

Q Great magazine, it's made moving from Windows to Linux painless for me so far, thanks. I am having one problem, however: I'm using *Mandrake 7.2* on my laptop and I'm trying to learn C++. Using gcc (2.95.3), the code I've been writing has been compiling fine, but BASH gives me a 'command not found' error when I try to run anything. I thought it might have something to do with the difference between a.out and ELF, but according to the gcc HOWTO, both my system and the compiler are dealing in ELF. (Interestingly, if I don't specify an output file it creates one called "a.out".)

Any ideas what I could be doing wrong?

A If you don't specify an output file, gcc, along with many other compilers, will just call the executable a.out. there's no problem there. I think what you are doing is failing to use BASH properly. There are two ways you can execute a command or program using BASH and most other shells. Firstly, you can place the executable file in the **PATH** that BASH checks for commands. The **PATH** is normally defined in the .bashrc file you'll find in your home directory, and is stored as an environment variable. Type: **echo \$PATH** to see the current path on your system.

The other way to execute a program is to give its path. When there is a path in front of the command, BASH knows it isn't expecting to find it in the paths it normally checks for commands. You

have to do this, even if you are in the same directory as the file you want to execute. In that case you would use the . which is interpreted as the current directory: **./a.out**.

This should solve your problem.

Modem solution

Q I notice questions about connecting to the Internet under Linux, and the answers seem to be rather complex for what should be a simple task. Can I offer the method I use? This is all done as root:

1. Ensure that the *wvdialpackage* is installed. I have just installed Red Hat Linux 7, which includes *wvdial-1.41-10.i386.rpm* therein, so I would presume that is the current version.

2. Run the following command:

wvdialconf /etc/wvdial.conf

This will auto-detect your modem and insert the relevant details in the file stated. On my system, it created the following file:

[Dialer Defaults]

Modem = /dev/ttyS1

Baud = 115200

Init1 = ATZ

Init2 = ATQ0 V1 E1 S0=0 &C1 &D2 S11=55 +FCLASS=0

; Phone = Number

; Username = Username

; Password = Password

3. Edit */etc/wvdial.conf* and delete the three comment lines at the bottom, then insert a separate block of lines at the bottom with the following format, separated from the previous blocks of lines by a blank line:

[Dialer FreeISP]

Phone = 08081 570 980

Username = ThisIsMe

Password = Secret

The word in the square brackets after "Dialer" should be an easy to remember mnemonic for the ISP in question, and the phone, user name and password should obviously be the appropriate ones for the ISP you're using.

The phone number shown is from the 'Reserved for Drama Use' section on OfTel's website, and is



Juggling drives can make dual booting easy, but beware of the picky Windows file system!

guaranteed never to be issued. Refer to <http://www.oftel.gov.uk/numbers/drama599.htm> for details.

4. Save the file.

When you wish to go online, the command is **wvdial FreeISP** with the 'FreeISP' replaced by the mnemonic chosen.

Riley Williams
Aberdeen

A Thanks for your advice, which I'm sure will be of help to RH7 users. Our previous instructions have concentrated on packages such as *kpp* (provided with KDE), as these tools are more widely used.

More Mandrake malarky

Q I have installed Mandrake 7.2, as delivered on the CD with the Christmas 2000 issue, onto my computer (a Compaq with a 6Gb hard drive, 128Mb RAM and an AMD processor) and followed the instructions as best I could.

This has given me something of a dilemma. The machine, along with a laptop is used as a business tool and also as a domestic machine for the family. The implementation of Linux has introduced a load problem in that during the sequence the user is presented with a short menu of operating systems to choose from: Linux, Fail Safe, Windows or Floppy.

This is fine except that the menu times out after a few moments and defaults to Linux – which is not very practical when one's mind is elsewhere assuming that the load will offer the usual Windows interface.

It will be a while before I can totally adopt Linux and this niggling and inconvenient situation will cause some bother. Is it possible to rearrange the menu so that the default becomes Windows?

Some observations:

1) I use two mice – a wired in or infra red remote. Each has a wheel. None of the options for inclusion during the configuration section of the load procedure would accept these devices. I got the USB mouse off the laptop and it worked OK.

2) The method of getting online (I use Netscape 4.6) failed completely.

Perhaps I did something wrong but it seems awkward.

3) Is it possible to drag files from Windows into the Linux utilities (for example, my business files such as letters and spreadsheets)?

4) Is there a good general reference book I can buy about Linux?

Thank you for a well structured and informative magazine.

A The problem you have with booting is easily solved. First, you'll have to be logged in as root. Then open the file `/boot/grub/menu.lst` with a text editor – it doesn't matter which one, KWrite will do fine. Near the top of the file you'll see a line which reads line

```
default 0
```

This tells the boot menu which is the default option in the menu. In your case, you want to change the '0' to a '2' to make Windows the default (the count begins at 0, so the third entry is identified as 2). Save the changes, reboot your machine and it should default to loading Windows.

As for your other points:

1) You need to supply more information here for us to help. What exactly are the types and models of the mouses you tried.

2) Again we could do with more information. What type of modem do have? If it is an internal PCI model (or Winmodem) then it is probably not supported by Linux.

3) Yes. First you have to tell Linux to access or mount the Windows partition on your drive. Again you need to be logged in as root.

When you mount a drive, you need to mount it somewhere within your existing filesystem. Typical places are `/cdrom` or `/mnt/cdrom` for the CD-ROM drive `/floppy` or `/mnt/floppy` for the floppy drive, etc. For your Windows partition, we'll create a directory


called `/mnt/windows` for it to get mounted on. First check whether this directory exists already. Enter **ls /mnt**. This will list the contents of the directory `/mnt`. If there's not a directory in there called `windows` then create it with **mkdir /mnt/windows**.

Okay, you're ready to mount now. Enter **mount /dev/hda1 /mnt/window -t vfat**.

This tells Linux to mount the first partition of your first IDE drive on the directory `/mnt/windows` and that the filesystem used on this partition is 'vfat', which is what Linux uses to read Windows FAT filesystems.

You should now be able to browse your windows drive in the file manager and drag-and-drop files on to your Linux drive. Whether you can open those files in any meaningful way when you get them there is another story.

Instead of typing in that long mount line every time you need to access the Windows partition, edit the file `/etc/fstab` and add the line



SUBMISSION ADVICE!

We welcome all of your technical queries, and believe me, all of them will be read. We want to give you the best help possible, so if you are writing in, please try to follow these simple guidelines.

- Give as much relevant information as you can – at the basic minimum, it usually helps to know what flavour of Linux you are using. Other problems may be easier to solve if we know specific information on hardware, software versions, etc.
- Be concise – if you tell us all about your life so far, the price of a frozen pizza and how strange the weather has been before you get to your problem, chances are we'll have fallen asleep.


Please remember that under most circumstances we won't be able to reply to your query personally.

You can send your questions in various ways.

The postal address is: Linux Answers, **LINUX Format**, 30 Monmouth Street, Bath BA1 2BW, or you can email us at linuxanswers@futurenet.co.uk

You can also ask online by visiting our online forums. Just go to this address: www.linuxformat.co.uk/forums

/dev/hda1 /mnt/windows vfat rw defaults,user,noauto to the bottom and save it. You'll then be able to mount the drive with **mount /mnt/windows.**

4) There are hundreds of Linux books out there. For the beginner, we quite liked Linux in Easy Steps (ISBN1-84078-072-X). 

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 worldwide, 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 Keron

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

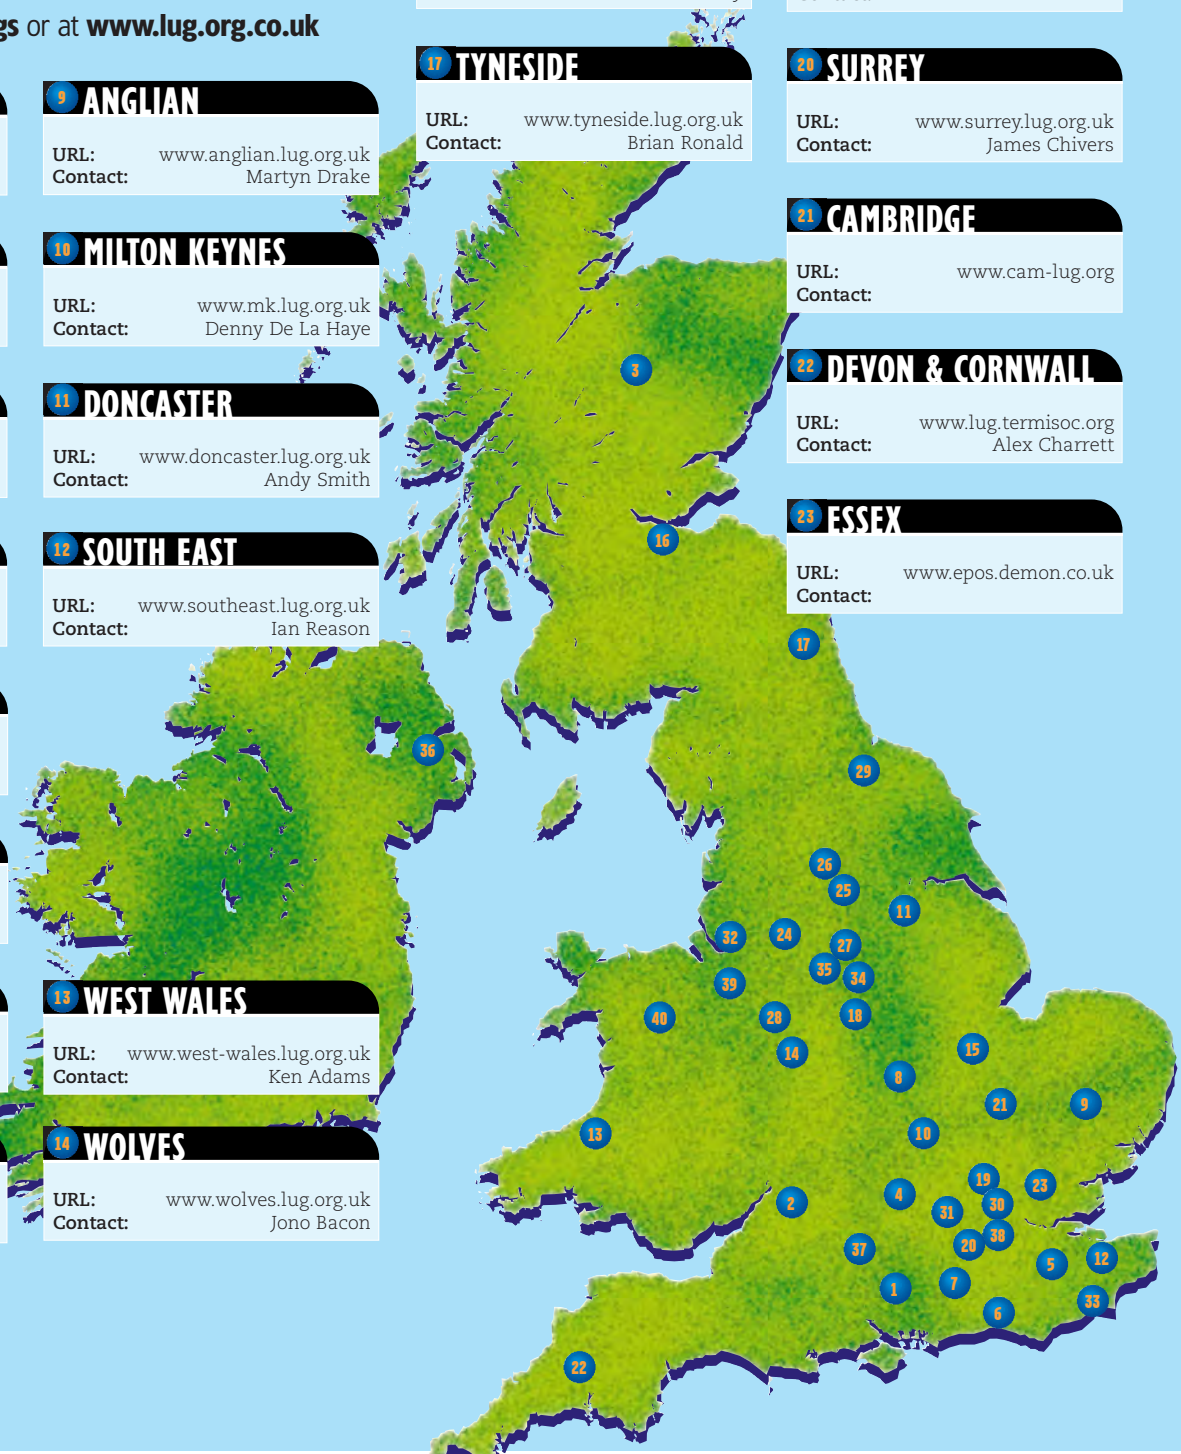
19 GREATER LONDON
URL: gllug.linux.co.uk
Contact:

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

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

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

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



USER GROUPS

24 MANCHESTER

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

29 NORTH-EAST

URL: www.shofar.freemove.co.uk/NELUG

34 CHESTERFIELD

Email: info@spirelug.org.uk
Contact: Paul Sims

39 CHESHIRE

E-mail: enquiry@sc.lug.org.uk
Contact: Richard Smedley

25 WEST YORKSHIRE

URL: www.scs.leeds.ac.uk/wylug
Contact:

30 LONDON

URL: www.lonix.org.uk
Contact:

35 SOUTH DERBYSHIRE

URL: www.sderby.lug.org.uk/
Contact: Dominic Knight

40 NORTH WALES

URL: www.northwales.lug.org.uk
Contact: Dr. Tim Traveres

26 WEST YORKSHIRE

URL: www.wylug.lug.org.uk
Contact: Nigel Metherringham

31 THAMES VALLEY

URL: www.sclug.org.uk
Contact: Nick Lambert

36 BELFAST (BLUG)

URL: www.belfastlinux.org
Contact: dannywalk@ntlworld.com

27 SHEFFIELD

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

32 LIVERPOOL OpenSource

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

37 WILTSHIRE

URL: www.wiltshire.lug.org.uk
Contact: Jason Rudgard

28 STAFFORDSHIRE

URL: linux.ukweb.nu
Contact:

33 DEAL AMIGA CLUB

Email: superhighwayman@hotmail.com
Contact: John Worthington

38 SOUTH LONDON

URL: www.b-Lug.org
Contact: Dr. Tim Traveres

WORLDWIDE LUGS

AFRICA

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

AUSTRALIA

ADELAIDE LUG
Url: www.linuxsa.org.au
Contact: Matthew Tippet
Email: mtippet@anu.edu.au

MELBOURNE, VICTORIA
Url: www.luv.asn.au
Contact:
luv-committee@luv.asn.au

PERTH
Url: plug.linux.org.au

EUROPE

EIRE
Url: www.linux.ie
Contact: Ken Guest
Email: root@linux.ie

URL: www.dilu.org
Contact: Joe Lennon
Email: glossary@dilu.org

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

INDIA

Url: www.river-valley.com/tux/index.html/
Contact: K. Anilkumar
Email: anil@river-valley.com

NORTH AMERICA

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

CLARKSVILLE, TN
Url: <http://www.cllug.org>
IRC: [#Linux](irc://irc.midsouth.net)
Email: tux@cllug.org

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.br lug.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

LUG ORGANISERS!

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

Name of LUG: _____

Location: _____

Contact Name: _____

Website address: _____

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

SOUTH AMERICA

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

BUENOS AIRES
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

CREATE A DATABASE IN EXCEL ⚡ 20 TIPS WINDOWS MEDIA PLAYER 7 ⚡ SIX PAINT PACKAGES ON TEST

PC Answers[®]

TIPS ✓ PROJECTS ✓ HELP FOR YOUR PC



Print faster, better

Essential tips to help you get more from your inkjet



Outlook Express

FAQ: fix error messages and solve other email problems



Your own programs

Create your own custom tools with our CD and new series

PC Power Tips

✓ Speed up ✓ Protect
✓ Customise ✓ Control



⚡ Hit every search engine with one click!
⚡ All you need for future-proof protection from email viruses



132 pages packed with hints, tips, tutorials and technical help!

On sale 22nd January

Order your copy today or call ☎ 01458 271127 to subscribe