

EDITORIAL

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

Art Editor **Chris Crookes**
chris.crookes@futurenet.co.uk

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

Production Editor **Andy Channelle**
andrew.channelle@futurenet.co.uk

Editorial Contributors

Rob Fenwick, Dave Goulson, Mike Saunders,
Steve Heaven, Jonathan Kent, Charlie Stross,
Kate Hadley, Chris Howells.

ART CONTRIBUTORS

Photography **Rick Buettner**, **Kath Lane-Simms**,
Louise Parker, Superstock Images,
Illustration **Gary Kempston**, **Garry Parsons**, **Tony Sigley**

ADVERTISING

Group Advertising Manager **Mike Roberts**
Deputy Advertising Manager
Michelle Blackwell: 020 7317 2602
Display Sales execs
Lorien 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**, **Stu Hobbs**, **Melissa Stapleton**,
Chris 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 **Ben Tisdall**
Assistant 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

Chris Power: 01225 442244
Fax: 01225 732384 cpower@futurenet.co.uk

Contact Details

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

Subscriptions, back-issues and Mail Order

Phone: 01458 271178. See page 80
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.

Printed in the UK by Midway Clark (Holt) and © Future Publishing Ltd 2000
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 • MILAN • MUNICH • NEW YORK • PARIS
ROTTERDAM • SAN FRANCISCO • TOKYO • WROCLAW

WELCOME

Welcome to the very first issue of *Linux Format* in 2001. Sadly we have not seen any giant monoliths springing up in the office, nor has anyone been caught brandishing a bone as a weapon, but there's been plenty of Linux activity going on.

Though it still may seem a little premature, we have a comprehensive look at the new kernel this issue. The kernel is the very heart of Linux, and this version has been eagerly awaited for over a year. For many, this is probably the most important new version of the kernel since it started, because it adds features which will make Linux truly capable of Enterprise-level computing, with support for symmetrical multi-processing, amongst other nice features. Desktop users haven't been left out either – for the full story of what we can expect, plus some ideas of where kernel development is headed, take a look at our feature starting on page 46.

One of the best things about Linux is that it's all about choice. You can choose whether you want to use *KDE* or *GNOME* for example. You can also choose which web server you want to run. Many people may not have heard of *Roxen*, and if you're one of them, you should definitely check out our guide to this great software on page 52.

There's really too much in this issue to explain here, so head on in for the best Linux treat of 2001 so far...

Nick Veitch
EDITOR



Chris Crookes
Art Editor

Mad Chris's favourite musical track of the moment is *Dusty's You don't have to say you love me*



Richard Drummond
Reviews Editor

Rich's favourite track this month must have been *Blondie's early 1980 track, Sound-a-sleep*

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

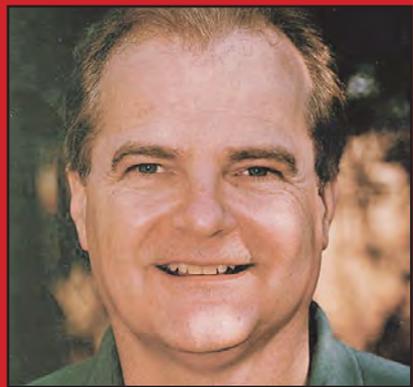
CONTENTS

Linux Format, January 2001 – your new year of the open source revolution starts here!

Interview

Craig Hunt, Linux author and network guru talks to Nick Veitch **p14**

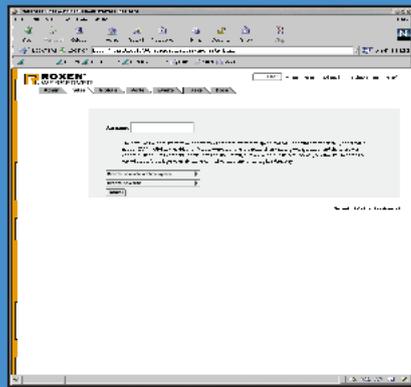
We caught up with long-term Unix expert and author Craig Hunt, and probed him for his views on everything Linux.



Roxen

Does Apache finally have some worthy competition? Check out Roxen **p52**

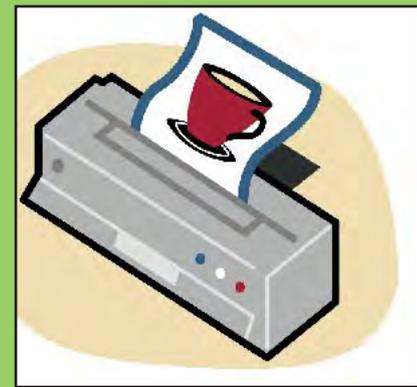
The Roxen webserver may be one of Linux's best kept secrets – but no longer! Find out what this great server has to offer here.



What on Earth

Problems with your printer – you need CUPS. Linux printing that works **p62**

The CUPS software is included on our CD this issue, but find out what it does and why you want it here...



REGULARS

NEWS	8
GNOME and KDE get backing, Euro patents postponed, SuSE goes PPC again and more...	
SUBS	84
Subscribe to LINUX Format and save money, never mind shoe leather. All the details are here.	
COVERDISC PAGES	90
Your guide to another CD packed full of useful Linux tools, source, demos and docs!	
NEXT MONTH	98
Here's a preview of what we'll be covering next issue, if you're still thirsty for more!	
LETTERS	58
<i>Your views and comments.</i>	
LINUX USER GROUPS?	94
<i>You're never alone with a LUG apparently, and there are plenty of them...</i>	

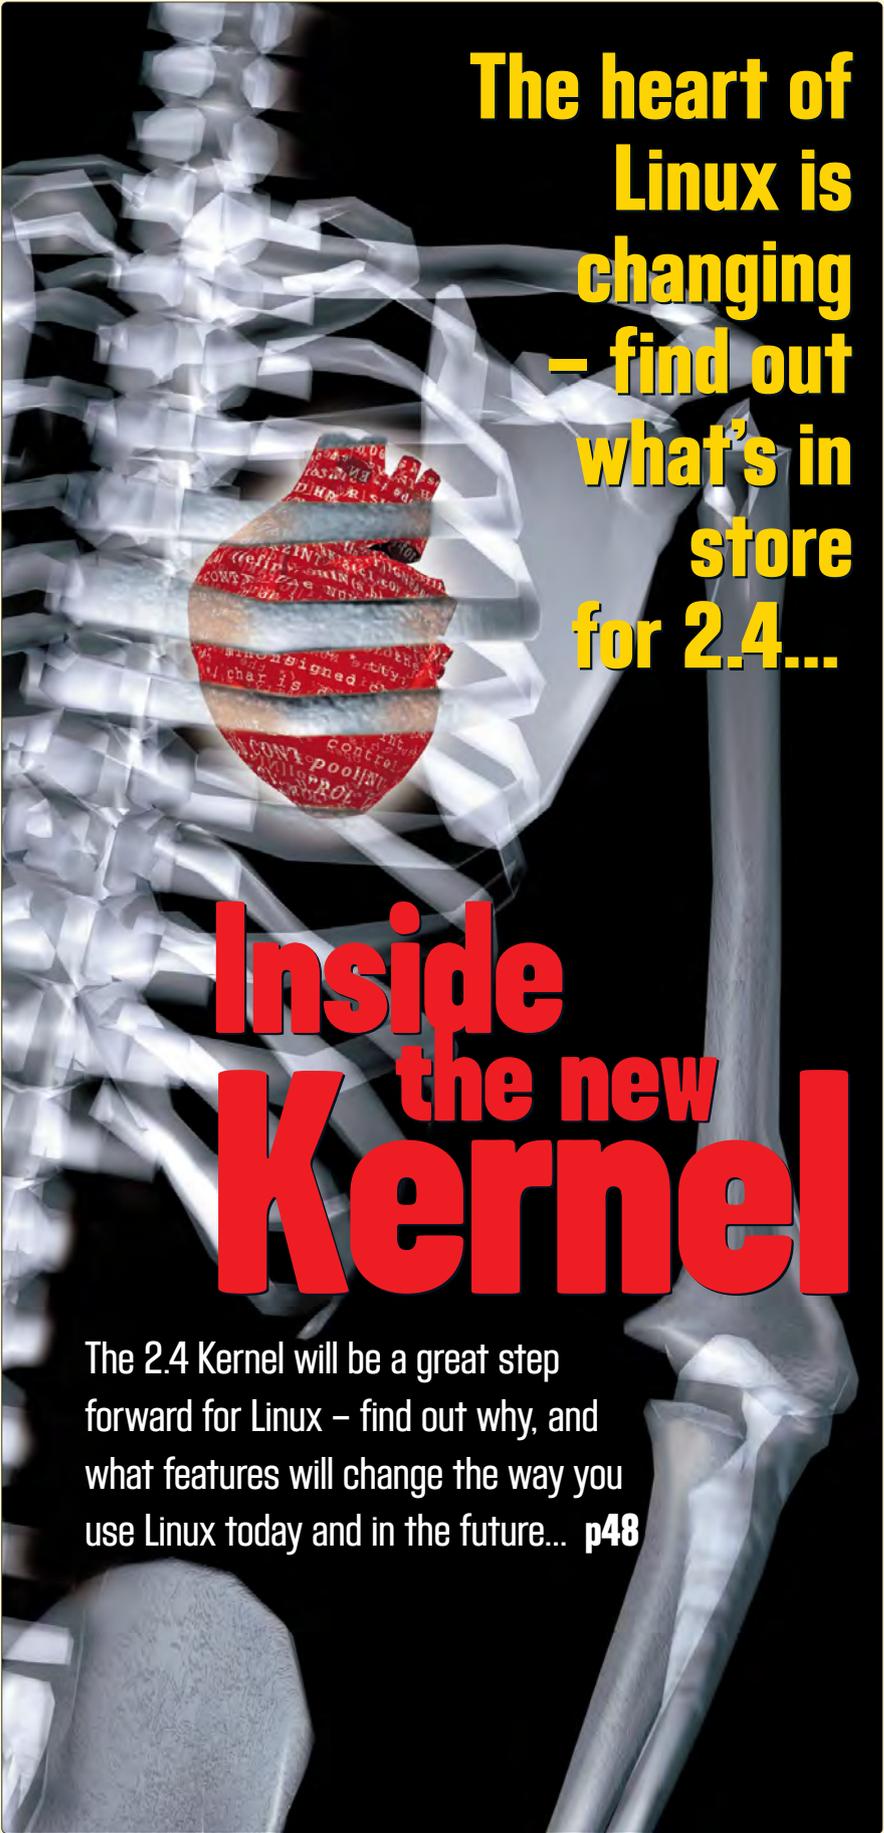
REVIEWS

MANDRAKE 7.2	18
Mandrake aim to stay ahead of the pack	
OMNIS STUDIO	20
Web-based application development on test.	
TIME NAVIGATOR	22
Easy to use, extensive backups.	
OPERA	24
A preview of a browser you might want to pay for!	
SHOGO	26
Robots have run amok on the factory floor...	
HEAVY GEAR II	28
More robot mayhem – but whose robots are best?	
BOOK REVIEWS	30
DNS, Internet shopping and more covered this month	
IRC ROUNDUP	32
The best and worst of Linux IRC clients exposed in our extensive roundup	
HOT PICKS	40
The latest and greatest open source software, rounded up and filtered for your edification!	

TUTORIALS

PERL PRIMER	68
One of the most interesting languages available to Linux programmers, explored in part one of this series.	
JAVA	73
Exceptions, input, output and the classes that control them explored in this month's episode.	
PILOT	77
Synchronise your Palm Pilot to your Linux desktop machine the easy way – follow our guide!	
NFS	80
Simple Unix networking made, er, simple. NFS explored and explained this issue.	
LINUX ANSWERS	86
<i>Our experts will solve all your Linux problems from PPP and disobedient printers, to multiple installs and more...</i>	



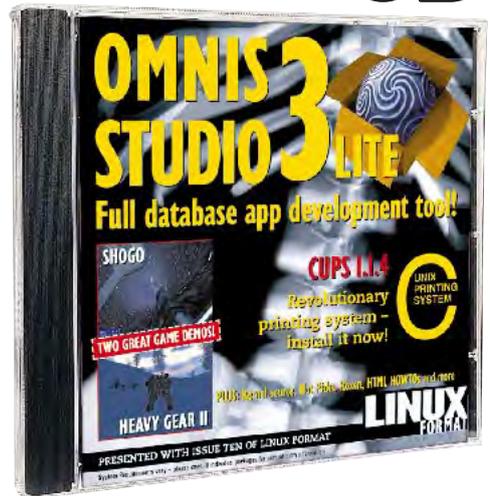


The heart of Linux is changing – find out what's in store for 2.4...

Inside the new Kernel

The 2.4 Kernel will be a great step forward for Linux – find out why, and what features will change the way you use Linux today and in the future... **p48**

LINUX FORMAT CD



FULL APPLICATION DEVELOPMENT SYSTEM AND MORE ON YOUR CD!

Find the latest web servers, game demos, kernel source, up-to-date documentation, printing software and much more on this issue's CD!



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

Please read the coverdisc instructions on page 90 before installing from the CD!

NEWS

Another busy month: KDE get their act together, Europe gets a patent reprieve, Linux gets a new MP3 player, Framemaker gets dropped and Crusoe gets bugged...

ANYTHING GNU CAN DO...?

Hot on the heels of the creation of The GNOME Foundation comes an announcement of The KDE League.

The developers of *KDE* have banded together with a number of big names from the PC and Linux world to 'promote the distribution and development of *KDE*'. Founding members of the board include Borland, IBM and Hewlett-Packard and distribution vendors such as SuSE, Mandrakesoft, Corel, TurboLinux and Caldera.

“We are simply trying to reach a different kind of user that does not read geek sites.” – Chris Schlaeger

However, in contrast to the GNOME Foundation – created with the aim of ‘steering’ development of the desktop environment – the KDE League will not be directly involved in the *KDE* codebase. Instead, the various partners will provide ‘financial, moral and promotional’ support with three principle goals: to ensure *KDE* remains a desktop standard on PCs; to help *KDE* compete with the prevalent proprietary desktops; and to encourage third party application development for the environment.

Matthias Ettrich, founder of *KDE*, said: “The creation of the KDE League marks a vital step forward.

With the support of our corporate partners, we can work together to ensure that *KDE* gains wider recognition, a greater number of applications and increased functionality, while maintaining the open development model.”

In tandem with the creation of the League, were a few high profile announcements including IBM revealing that it was working to bring *ViaVoice* – the company’s voice recognition system – to *KDE*.

All this is a very long way from *KDE* Developer Kurt Granroth’s assertion that a GNOME Foundation-style organisation “flies

square in the face of everything the *KDE* project stands for.” However, the organisation claim that the companies involved in the League won’t have any extra influence in the development style of the project in any way: “The League exists solely for promotion and education.”

In an open letter to the developer community, Chris Schlaeger said that *KDE* was currently used on 70% of Linux desktops and rather than trying to grab the other 30% of Linux users, the League should target the “95% of desktop users [using other OSs] rather than compete with our friends on the *GNOME* project. Just

converting 5% of Windows users will get us more *KDE* users than converting all *GNOME* users.”

He also stressed the League’s Public Relations-only role, saying: “We are simply trying to reach a different kind of user that does not read geek sites or mailing lists.”

In response to the charge that the KDE project was merely copying the GNOME Foundation, Schlaeger writes: “When the GNOME Foundation was announced it was a bit of a wakeup call. It was not the Foundation itself that bothered me, but the fact that many people believed KDE was doomed because it had no commercial backing.”

We are The League

Current members of the KDE League, as announced at Comdex, are Borland, Compaq, Corel, Fujitsu Siemens, Hewlett-Packard, IBM, KDE.com, Klaralvdalens Datakonsult, theKompany.com, Mandrakesoft, SuSE, Trolltech and TurboLinux.

Founding Fathers

And the membership of the GNOME Foundation includes Compaq, Debian, Eazel, Free Software Foundation, Gnumatic, Helix Code, Henzai, Hewlett-Packard, IBM, Mandrakesoft, Object Management Group, Red Hat, Sun Microsystems, TurboLinux and VA Linux.



KDE have secured the backing of some big names in their bid to become the default desktop.

COMMENT

A lot of media mileage has been given to the fact that KDE and GNOME are competing desktop systems. In a sense they are competing, but in a positive way. The developers of KDE can look at GNOME and say, “Hey, that’s a great idea. How do we implement that?” or vice-versa. It’s not competition in negative way, since the user, having decided to try either KDE or GNOME is not then locked into that choice. You can run both systems on the same machine, even at the same time if you want to. And just because you use GNOME’s desktop doesn’t mean you have to use GNOME’s browser - *Konqueror* will work just as well. Conversely, *Nautlius* will run on a KDE desktop. This is possible through the power of open source development.



NO SOFTWARE PATENTS FOR EUROPE... YET

The long-running battle to prevent the European Parliament from extending patent laws to cover software achieved a partial victory when most of the member-states (the exceptions being Austria, Lichtenstein and Switzerland) agreed to put off the move.

The decision means that computer programs will continue to be exempted from patents at least until the European Commission completes an extensive public consultation on the issue. National governments will also be urged to participate in the debate.

However, the battle is not yet over. "We are still very far from a decision to ban software patents in Europe," said Stefane Fermigier, a member of EuroLinux, an organisation set up to lobby the commission.

US law has made provision for the patenting of 'elementary software processes' – despite



widespread opposition from developers and publishers – and Europe is under intense pressure to harmonise patent law to facilitate global recognition of patents. The organisation Freepatents claim that no proper study on the economic

The EU has ruled against software patents for now, but you could still visit www.eurolinux.org to sign their petition.

impact of software patents has been made in Europe and that legislation, if it's not carefully designed, would 'create a lot of judicial uncertainty, playing havoc with the whole European independent software industry, obstructing competition and, in the mid term, slowing down innovation and investment'

GREAT SOFTWARE PATENTS OF OUR TIME

British Telecom hit the headlines recently by claiming to hold a patent on the hypertext link, and there are plenty more where that came from...

NUMBER: WO0054187

PURPOSE: A 'universal media player' that has visual representations of various media (i.e icons) which, when clicked on, are visually loaded into the player.

NUMBER: WO0062265

PURPOSE: A method of securely distributing music and video across



a public network using a client/server architecture.

These may seem small and inconsequential, but if you're a software developer who doesn't have the time to wade through 150 million pages of documentation before adding a feature to your media player, you could easily fall foul of the law. And end up in court.

There is also the prospect of fundamental or even 'standard' processes being patented, such as BT's attempt at claiming the hyperlink or Microsoft allegedly trying to patent portions of the W3C standard, which will have repercussions for us all.

European GnuS

A group of individuals from Germany have banded together to form a European sister organisation of the Free Software Foundation.

The organisation, affiliated to the original FSF in America, will provide a political mouthpiece for free software and the GNU project in Europe. In the launch announcement, Georg Greve said: "It has become evident that the FSF in the US, being in a totally different time and culture zone, can hardly keep in touch with the developments in European Free Software."

If you'd like to help them become a truly European group, contact: team@fsfeurope.org.

News in brief...

Plex86, the 'virtualisation' software has hit another landmark in its development. Kevin Lawton, one of the leading developers of application, recently added another notch to his OS bedpost by getting **Windows 95** running concurrently with **Mandrake** on his Linux box. So if WINE's fallen over a few times too many and the cost of VMware (which has just increased up to a wallet busting \$299) seems an extravagance – and you need Windows – you now have a third option. Plex86 can now run Win 95, MSDOS, FreeDOS and Linux as guest systems with support intended for Win 98 and NT.

FreeBSD 4.2 has been released and is available for download from www.freebsd.org. The OS is available at present for **i386** and **alpha** platforms and fixes a number of bugs in the previous release. If you can't justify the inevitably huge download, a four-CD pack should be available as you read this from www.freebsdsmall.com.

NetBSD have announced release 1.4.3 of their OS. The project site (www.netbsd.org) says this release has resolved some stability issues and fixed over 70 bugs reported by users.

The future is now... or it was back in November when with world's most technologically advanced **robots** converged on Yokohama, Japan for **Robodex 2000**. The biggest hit of the show was Sony's **SDR-3** (Sony Dream Robot 3); a two legged prototype with rudimentary voice recognition, a 180,000 pixel CCD 'eye' and the ability to kick a ball.

Macmillan Publishing are marketing **Linux For Windows**, a product based on Mandrake 7.2 and promising a simple three step install and no partitioning. The package is completed by **KOffice** and a series of multimedia tutorials on subjects such as installation, CD writing and file management.

Compaq are expanding their Linux operations by announcing they will now offer pre-loaded Linux on a range of servers and desktop →

News in brief...

→ computers. "Compaq is delivering the most robust Linux solutions, meeting our customers' need for greater flexibility and expanded choice in developing their e-business platforms," said a spokesman.

Microsoft's latest announcement is causing a little gnashing of teeth as it seems **Whistler** (Windows' eventual replacement) *might* be set up to only run certified (by Microsoft, natch) code. Does this mean shareware developers will have to submit to **Bill Gates'** personal scrutiny? And will Sysadmins be able to turn it off? Something for the DoJ to think about as Bill begins to appeal against the breakup of his monopoly, er company...

Will **distributed computing** be the Next Big Thing? Compaq, **Hewlett-Packard** and **SGI** have teamed up with **Platform Computers** to form the New Productivity Initiative to harness the untapped power in today's computer systems. The most high profile use of distributed computing so far is the **SETI@home** project which searches for alien intelligence while you take your lunch break, but new software should allow for even the spaces between the double-click of your mouse to be utilised. Your computer never need be idle again...

The **Albuquerque Journal** is reporting that **IBM** have supplied the University of New Mexico with a **Linux Cluster** that they're claiming is 'the fastest computer ever built'. This super cluster will be used by the US Airforce to process images from a telescope on top of Mount Haleakala in Hawaii to identify objects in orbit around Earth.

In the past, images from the →
(Continued on page 12)



The International Space Station

LINUX IN ACTION UPDATE

Eyes-free surfing for Linux

Emacspeak inc. recently announced the worldwide release of version 13.0 of their revolutionary audio desktop software.

Emacspeak is a fully functioning environment that brings completely eyes-free access to Linux, allowing seamless browsing of local and remote information with a consistent, integrated UI. With the addition of the freely available **IBM ViaVoice** TTS speech synthesis engine, Linux has become the first no-cost internet solution for blind and visually impaired users.

The key developments in this version of the software include the ability to use sound themes for auditory icons, streaming media through *FreeAmp* and *TRPlayer*, single click access to mp3 playlists, XML and SGML compliance and a simple speech-enabled calculator.

Emacspeak is included voluntarily with most Linux distros and the latest version can be downloaded from <http://emacspeak.sourceforge.net> or <ftp://ftp.cs.cornell.edu/pub/raman/emacspeak/>.



Get a penguin in your cube

Intel users have had the opportunity to discover the pros and cons of SuSE 7 for a few months but their Mac owning friends were excluded from the party. Until now.

SuSE 7.0PowerPC Edition is a five CD set featuring over 1000 applications including all the usual



suspects including *GIMP*, *KOffice* and a suite of HTML editors. One surprising addition is *Mac on Linux*, an application which – SuSE claim – will allow users to run the MacOS from within Linux, removing the need to run a dual boot system.

The familiar YaST2 tool has been tweaked for easy installation, SaX2 has been added for configuration of graphics cards and support has been added for IBM's pSeries processors, making this a good distro choice for IBM RS/6000 users.

An uncertain future for Corel Linux?

After months of speculation following a cash injection from Microsoft, Corel clarified their position with regard to the Linux market by suggesting that changes were essential if the company was to survive and thrive.

CEO Derek Burney said that the company was now seriously

considering its options, which may involve selling or expanding their Linux operation. The announcement caused a sharp dip in the company's share price.

Inevitably, Microsoft's critics were quick to suggest that Corel's move away from the Linux

platform and onto .NET applications would be quite convenient...



FAX MAGIC

Merlin Software used November's Comdex show to unveil *Communicado Fax*, a commercial application which should bring print to fax capability to the Linux masses.

Fax to print is nothing new, but this is the first application to let any user on a Linux network (using any operating system) to print from a range of office productivity software including *StarOffice*, *WordPerfect* and *Microsoft Office*.

Based on the popular *HotWire Fax*, *Communicado* has all the expected facilities, including broad hardware and application support, a built in viewer with thumbnail and coverage facilities, detailed usage reporting and support for multiple in and outbound lines.

If all this network functionality seems like overkill, and you just want to fax the occasional letter from *StarOffice*, the personal edition of *Communicado* is also available from: (www.merlinsofttech.com).

For an open source alternative, check out Gefax at <http://gefex.sourceforge.net/>

Juke of URL

Musicmatch Jukebox, which already has a large following on both Mac and Windows finally gets a Linux release, and is available for download right now.



This fully skinnable MP3 player is capable of ripping and encoding MP3 files and features a very useable graphic equaliser, a very comprehensive library system for organising your music and support for visualisations. *Musicmatch* also has a jukebox facility which can be used to burn CDs from your playlist.

Musicmatch Inc called on the experiences of CodeWeavers, a company specialising on porting applications from Windows to Linux. Like many other recent products from companies such as Deneba and Corel, *Musicmatch* utilised *WINE* to speed up development of the Linux version, a fact which,

according to Dennis Mudd, President and CEO of Musicmatch inc "makes it easier for the Linux and Windows version to evolve in sync."

Although some users will be put off by its use of *WINE*, *MusicMatch* is a very capable MP3 player which brings cataloguing, net radio and CD burning into one easy to use and great looking package. www.musicmatch.com

COMMENT

There are good and bad points about developers using *WINE* to port apps to Linux. On the plus side we end up with applications that would otherwise take too much time and money to develop using native code. It's true that people using *WINE* are not developing native applications, but then if they weren't using it, would they be developing anything for Linux anyway? *WINE* gives coders a simple(fish) way of opting in to Linux.

PDA for the masses

Empowerment through technology is the goal of a new program launched jointly by the MIT Media Lab and Harvard University. Digital Nations was launched in October to address the fundamental paradox of technology which meant that "the people, groups and nations that benefited most were those that were the best educated, most affluent and most powerful.

One of the key initiatives of Digital Nations is to bring down the high cost of accessing technology and a few interesting devices were demonstrated at the launch.

Top of the bill was Pengachu, a \$50 Linux-based handheld device developed by Rehmi Post, Wendy Ju and Matt Reynolds. This diminutive marvel promises built-in networking, MP3 audio and IP telephony, a 128x64 pixel (possibly) VGA screen and a long battery life.

The motivation behind the design was to bring the Internet to the masses, giving access to vital news sources, educational applications and peer-to-peer voice and data communications.

Whether Pengachu gets off the drawing board and into the hands of those who need it or not, the project does go to show that a cheap, accessible PDA is possible with Linux at its heart.

See <http://rehmi.www.media.mit.edu/~rehmi/pengachu> for more.

WEB WATCH...

Give your Linux box a makeover

The release of Netscape 6 has sent the LXF office skins mad. So, if you're tired of looking at the same old desktop, start tweaking and get



creative with your Linux box. www.themes.org is the granddaddy of all theming sites with sections for KDE, GNOME, AfterSTEP, Sawfish and all the other popular window managers. This is a great place to start if your interested in personalising your PC, you'll also find links to theming tutorials and pages where designers show off their latest creations.

If you're interested in Skinning Netscape's new baby, the first place



to start is the Themepark – <http://home.netscape.com/themes>. Here you'll find a few pre-made chromes, resources for skimmers and a download of Netscape's own Themebuilder tool. This application (currently in beta) is designed to allow the average user to begin skinning within about ten minutes.

Themebuilder should be great, but if you're not exactly DaVinci with a mouse, you might want to check out some pre-made backgrounds, button and bars.

Tiles are vital for backgrounds and, as they're smaller than a 1600x1200 desktop picture, are a great way to spare resources. There are thousands of sites on the web containing millions of tiles. Some of the best are at www.themes.org.

A more general resource is www.linuxartist.org which has a collection of useful resources for artists using Linux. There is a good news section and links to all the major apps.

Finally if you want your MP3 player to look as good as it sounds, check out Freeamp.org and download their latest player.

Once you're created your own masterpiece, why not submit it to www.skinz.org and let everyone else have a look.

DO YOU HAVE A HOT NEWS STORY?

The best way to let the Linux world know about it is to have it appear here!

Email us your news at linuxformat.co.uk

News in brief...

telescope were a little fuzzy, but this new computer will allow operators to examine these **UFOs** in greater detail, and decide whether they're missiles headed for the US or junk which could pose a threat to space craft like the **International Space Station**.

TurboLinux demonstrated the power of distributed computing at Comdex by using a small network and its **EnFuzion** software to crack five character WindowsNT passwords in just under a minute.

SGL are getting into the genome game with the installation of a Linux-based system for The **Bristol-Myers Squibb** Company. This bioinformatics system will use a Linux cluster from **Incyte** to research various genomics projects including DNA and protein sequence analysis, transcriptional profiling and proteomics.

The catchily named **www.completecomputercover.com** are now offering a range of **insurance** policies for computer users. David Dury, Sales Director said: "With CCC there are no difficult questions to answer, no minimum safety requirements and **no excess**, just a simple premium calculation and you have worldwide cover for your laptop."

Applix have moved to quash rumours that they are moving away from the Linux desktop market to concentrate on serverside applications. Spokeswoman Allison Antalek said that Applix were "committed to future releases of **Anywhere Desktop** (formerly *Applixware*)" and that the company would continue to provide the same "quality product and level of service to its existing customers" However, she also said that Applix were 'shifting the focus' onto server-centric and web-based applications...

A new company, **DSPsoft**, have announced a new '**Linux-like**' operating system for embedded devices such as set-top boxes and handheld computers. UnixCE will happily run in 340k of RAM and includes a FAT32, TCP/IP stack and a built in X-Windows API. A demo →

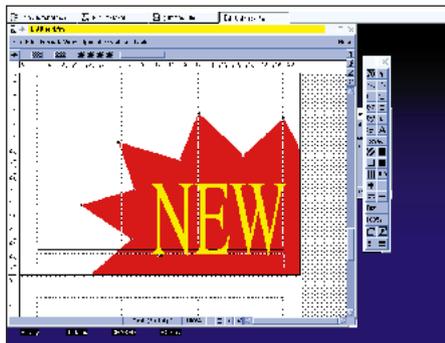
(continued next page column 4)

RIP FRAMEMAKER

Adobe have announced the suspension of their *Framemaker* for Linux beta program and said they will no longer be going ahead with a commercial release of their flagship document processor.

In an email to testers, Adobe said "The Framemaker product team

appreciates your feedback and participation in Adobe's investigations into the Linux market for Framemaker. From this study, Adobe has *determined* that it will not, at this time, release a commercial version of *Framemaker* or *Framemaker+SGML* on the Linux platform."



Betas of Framemaker will have expired by the time you read this.

What is not clear is whether the company took this decision because of an unfavourable reaction of beta testers to the application, the competition from 'free' software such as *StarOffice* and *KOffice*, or due to lack of faith in the Linux market.

Beta users were warned that the software was timed to expire on December 31 2000.

A bad day for chips

Two high-profile chip releases – Pentium 4 and Crusoe – were swiftly followed by even more high-profile withdrawals as 'flawed versions' and 'heat issues' respectively caused problems in machines.

Transmeta's shares tumbled by 15% after NEC recalled 'fewer than 300' systems running the chip, however, the company said the faulty chips came from a limited manufacturing batch.

In the US, Best Buy retail outlets withdrew Pentium 4 systems from sale after doubts were raised about 'excessive heat and inadequate performance'.

An offer you can't refuse?

Surf your way to a brand new free boxed distribution courtesy of Best Linux and Scotfree-uk.net

Scotfree-uk.net, the Linux friendly ISP, have joined forces with Best Linux to give *Linux Format* readers the chance to win one of 20 Best Linux box sets, complete with a Best Linux t-shirt, a cuddly penguin and a mouse mat. 200 runners up can also bag themselves a cool mousemat each.

So how do you get your filthy mitts on all this lovely booty? Simply visit www.scotfree-uk.net to enter – but make sure you get over there before 31 January 2001. And if you need a little help with the competition answers, a visit to www.bestlinux.net will sort you out.

Scotfree-uk offers users 0845 access to the Internet (or free access through Atlantic telecom) as well as 25Mb of webspace, an unlimited number of email addresses, free voice chat and webphone.

Best Linux are best known in

native Finland but have recently added English, Russian and Estonian versions of Linux to their product catalogue. Best Linux claim their latest release is easy to install and is ideal for home users. It ships with XFree86 4.0 and KDE2 as well as a suite of over 2000 applications including StarOffice and all the other usual suspects. The box also contains a games CD featuring emulations of DOS, Amiga, Atari ST, C64, Spectrum and Nintendo Gameboy, as well as a vast selection

of native Linux games. To complement the online help files and HOWTOs, you'll also get a comprehensive printed manual.



Get your hands on one of 20 BestLinux boxset for now!

NUMBER CRUNCHER

This month's totally meaningless statistics

6

Release version of Netscape's new browser, which has been receiving mixed reviews.

7.5

Weight (in pounds and ounces) of Linus Torvald's new daughter, Celeste Amanda.

45

Annual charge (in US\$) expected to be announced for users of MyMP3.com (and you can still only access your own mp3s!).

300

PCs recalled by NEC after the discovery of a serious flaw in Transmeta's Crusoe chip.

12,378

Open source projects hosted by Sourceforge.

90,000

Linux-based cash machines being installed by US based retailer Home Depot Inc.

110,277

Downloads of Ethereal (IRC client) from Freshmeat.net.

64,000

Maximum number of users that could connect to a Linux box before kernel version 2.4

1,000,000

Estimated number of user who will be left high and dry if Alta Vista abandons its email service.

560,000,000

Estimated number of text messages sent during the average month in the UK.

4.3,000,000,000

Maximum number of users able to connect to a Linux box running Kernel 2.4.

LINUX IN THE ARMY NOW

The latest convert to open source software is the New Zealand Army who have just installed Linux-based *Janus* software to train their grunts in the art of battle.

After some initial hiccups with the



Kiwis' newly acquired Dell servers, technicians from the USA's National Simulation Centre managed to get the kit up and running.

Major Julian Sewell, who is overseeing the installation, said that using the *Janus* software meant the army could press some old hardware into service, thus saving money.

Atteeeention! Do the New Zealand Army use the Born To Frag KDE Theme?

Eazel-Dell partnership

Computer manufacturer Dell have invested a significant amount of money in Eazel, the company responsible for *Nautilus* file manager (see Hotpicks), and say they will be bundling the application with their Linux laptops and desktop systems in the future.

As part of the deal, Dell are also expected to share in the revenue stream created by Eazel's subscription-based web services.

Eazel supremo Brian Croll believes the key to *Nautilus'* success is to make the software

'as widely available as is humanly possible'; and suggests that deals with other manufacturers could also be on the cards. "A distribution partner is key for us, but this is the first big one."



EMBEDDED LINUX NEWS

We tend to hear a lot about 'is Linux ready for the desktop?' But all over the world, developers and hardware manufacturers are toiling to stuff your favourite OS into a variety of useful devices.

FreePad is a new media terminal from Norway's Screen Media. Set for an early 2001 release, FreePad aims to offer Internet access, phone and answering services and uses DECT Technology (as used in cordless phones) to communicate with its base station.



The FreePad.

The idea, according to Screen Media boss Carl Henrick Janson, is to create "a new type of media channel for the kitchen and living room."

The thing that really sets FreePad apart from the competition is that it is, in essence, a complete Linux system so it should be able to run any Linux software.

Palm Palm Technology have unveiled the first Linux-based mobile phone. Designed for third generation mobile networks, the company claims the phone can handle a variety of multimedia functions and also playback

News in brief...

→ version for PCs is available at www.unixce.com. However, UnixCE is **proprietary** software requiring royalty payments to DSPsoft from developers who want to use it.

Linus Torvalds is reported to have signed a deal with HarperCollins to pen his **autobiography**. *Just For Fun: The Story of An Accidental Revolutionary*, will be co-authored by David Diamond and should be on the shelves in Spring 2001. One wag suggested he should have gone to Penguin (ho ho) for the deal, while another said this might be the reason for the delays in kernel 2.4. Lisa Berkovitz, associate publishing director, said the book would be "as much about Linux as it is about Linus"

A new survey (from the United States, of course) claims that programmers spend only **47 working days** a year actually developing applications. The rest of the time is spent **fixing bugs**, testing and coding for projects that later get cancelled (and drinking coffee).

French company **Bull** have invested in **Penguin Computing** with the intention that the two vendors will tackle the European Linux market together against stiff competition from the likes of IBM, HP and Dell.

The Agenda VR3, and transmission of audio files.

The Agenda VR3 is a portable PC running Linux which, its makers claim, represents an important step in the development of handheld computing. It also looks very, very, nice. British company Empeg - makers of an in-car mp3 player - have been bought out by SonicBlue, makers of the Rio brand of mp3 machines. Empeg claim the deal will allow them to focus their efforts on the fun stuff - product development.



MANDRAKE 7.2



Mandrake's latest release boasts a number of cool new tools. **Nick Veitch** picks up the seven CD set and puts their bragging to the test.

Web: www.linux-mandrake.com **Price:** £59.99

Of course, **Mandrake** has often been sold on its ease of use. But now other distributions are officially easy to install (including Red Hat), what Mandrake magic could be added to the installer to keep it ahead of the pack?

The first thing you'll probably notice is, thanks to the new kernel, Mandrake will allow you to use a USB mouse for the installation – very handy for those who wish to abandon PS/2 for good, though your mileage may vary depending on the specific device you use.

Disk partitioning is another area that often causes some terror and confusion in the hearts of those setting up dual boot systems, but *Diskdrake*, Mandrake's partitioning tool, is pretty reliable. Although it's always wise to have a back up just in case, we've had no problems with using it to resize Windows partitions.

Package selection is still a bit of a mixed bag. If you don't choose the packages individually (highly recommended, especially if disk space is tight), Mandrake will make a fairly good job for you, with some limitations. If you choose the 'development' option, the installer is more likely to select every library ever made and a compiler for every language you've used. **CUPS is now the default print manager in Mandrake.**

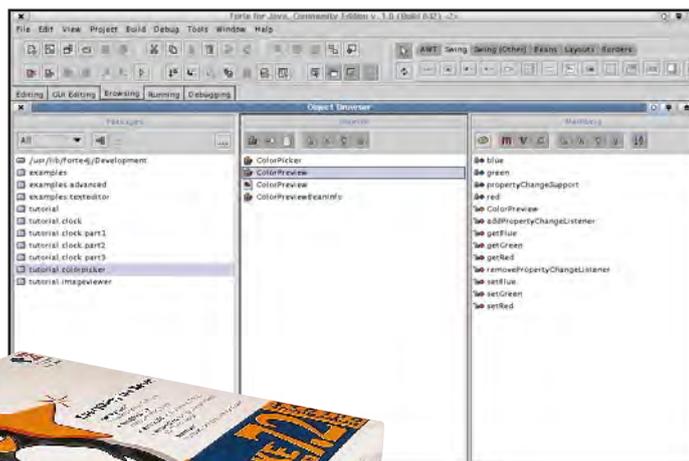


ever heard of – and some you haven't – before it decides to install some pretty fundamental X applications.

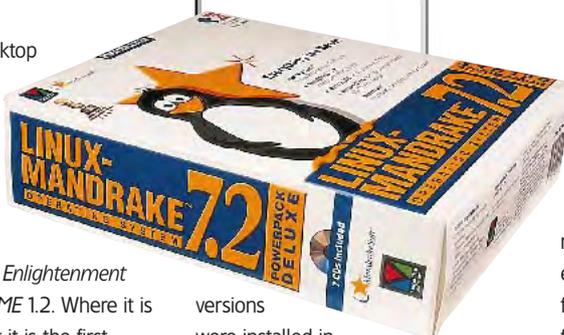
In terms of desktop environments and windows managers, Mandrake 7.2 still includes old favourites like *Windowmaker*, *xfce*, *Enlightenment* and of course *GNOME 1.2*. Where it is so far unique is that it is the first distro to totally abandon *KDE1* in favour of *KDE 2.0*, which was finally released shortly before the Mandrake 7.2 initial GPL disks (in fact, early GPL versions contained elements of the last beta version of *KDE2*).

It's all about choice

For many people this may be the first opportunity to try *KDE2* properly, and it is good to see that it is well supported in this release. If you are updating from a previous version of Mandrake and have already been playing with *KDE2*, it may be useful to note that Mandrake rolls *KDE* into the main filesystem (whereas *KDE2* test



The boxset contains a number of useful apps, including Sun's Forte.



versions were installed in `/opt` by default), so you may want to remove your test version first to avoid complications.

One of the things Mandrake has been most notable for in the past has been offering choice. Early on they adopted both *KDE* and *GNOME*, and more recently we have seen *Roxen* appear as an alternative – or addition – to *Apache* as a webserver. This time one of the most notable additions is the use of *CUPS* as the first choice print manager. *CUPS* is covered in more detail elsewhere in this issue (and is included on our coverdisc), but essentially, Mandrake sets this up in an easy to follow way,

much like the old *lpd* setup, and will even allow you to choose from the full range of printer drivers provided for in packages like *Gimp-Print*.

The joy of X

When it comes to X, Mandrake is pretty clever. *Xfree86 4* has been out for some time, but as we reported, there is little point in upgrading unless you have a supported card. The Mandrake installer attempts to figure out whether your card is supported, and then gives you the option of installing either *Xfree 4.01* or *3.6*. In addition, it will attempt to recognise cards for which 3D hardware acceleration is supported, and install the necessary drivers and libraries for this too – though it has to be said with mixed results. A test *ATI Rage*

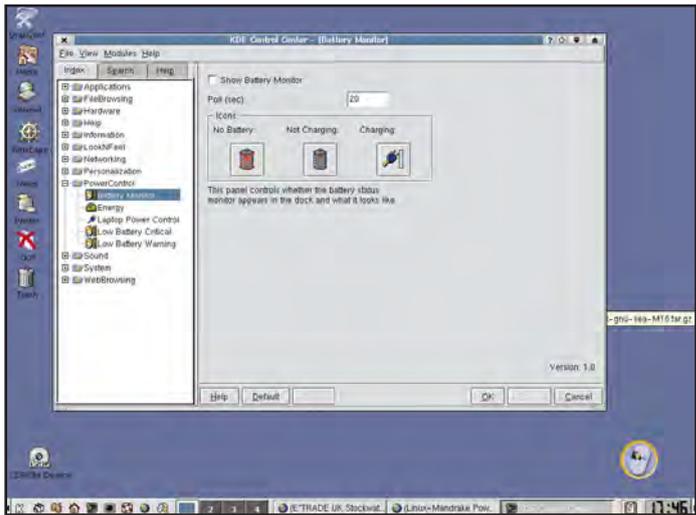
DOCUMENTATION

The Mandrake Deluxe edition features two printed manuals – an installation and users guide, and a Reference manual.

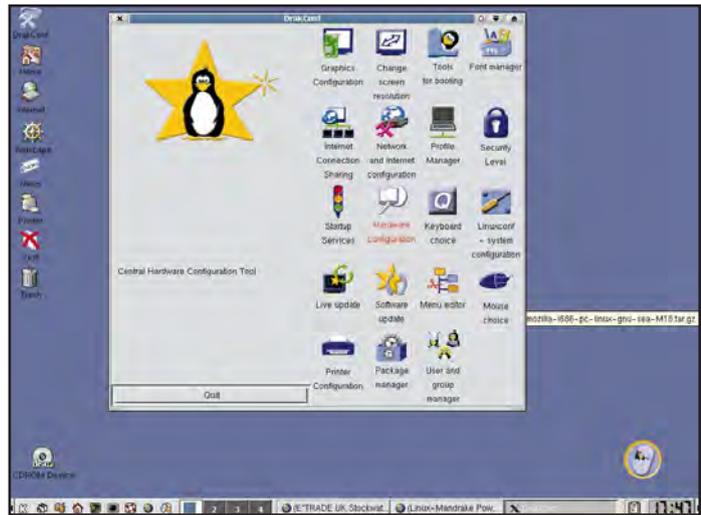
The installation manual, as you might expect, takes you through the install process and gives a fair amount of useful info on using the custom Mandrake tools. There's a useful glossary of

terms at the back, but the index is a little poor.

The 270 page reference manual is more thoroughly indexed and contains plenty of useful info on Linux in general, including some advanced administration topics. Both manuals are pretty accurate and easy to read, so top marks for that!



Mandrake ships with the final release of KDE2, but also gives you lots of choice when it comes to desktops.



Drakconf provides a natty graphical display for all your system administration tools and improves the Mandrake experience no end.

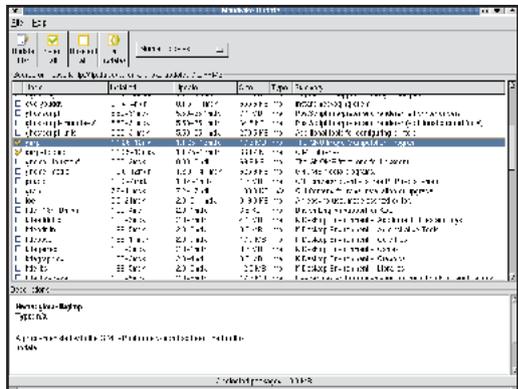
The auto-update is a good idea whose time will come, but not just yet.

128 worked fine, we had problems with a Riva TNT. Perhaps that's the difference between distros like Mandrake and,

for example, Debian – you may get less options in the latter, but you can be pretty sure it will always work.

One thing that might be worth taking on board for future versions is that I'm sure most people would prefer to configure X as early in the install as possible – after you've spent half an hour or whatever going through the install process, you don't want it hanging right at the end because your graphics card doesn't like being probed.

Another new feature is that Mandrake allows you to automatically log in as a specific user. Old-hands may gasp in horror at this – “the security risks!” – but the intention is obvious: to try and make the Linux experience as transparently easy as possible, and reproduce the kind of desktop environment that users are used to on other desktop systems such as MacOS and Windows. You could, obviously, set the default user as ‘nobody’ or a specially created guest user if you wanted. Or indeed,



not use this option at all. For people who are using a physically isolated machine, on their own, it's probably more of a blessing than a threat. In fact, if you want to point an insecure finger at Mandrake, you'd be better off complaining about the number of services which run by default on an installed system, without being properly explained beforehand.

Finally, like RedHat and some other distros, Mandrake now allows you to save a floppy disk containing this setup information, for automatic replication of the install, though sadly the floppy disk can't nip out and change the CDs over for you!

Keep is simple

So much for the install, what of the features? Linuxconf is usually the preferred choice of admin tool on any system, but Mandrake's own *Drakconf* is getting better all the time. Options for setting up networking are clearer and there is even a tool for setting up IP masquerading. The menu editor

COOL STUFF

As usual the extra disks in this boxed set include some software you might have difficulty getting hold of, and demo versions of commercial products. Featuring this time are some demos of professional apps like *Time Navigator* (reviewed this issue), *JBuilder* and *Code Forge*. One interesting demo version is the *CommunicatePro* message server.

There are also some full commercial products, like *EasyPrint Pro*, and IBMs *ViaVoice* package for Linux, and some stuff which isn't 'GPLed' like Sun's *JDK*, *Forte* and everyone's favourite, *StarOffice*.

There's a fair bit to play with, and a more useful selection of tools than the time limited demos that normally accompany such multiple disk sets.

can be used to easily edit the “top” portion of the menu of whatever desktop environment or window manager you are running (or even all of them at once if you like). This is more useful than it might actually sound, as you have a single interface for adding or removing applications onto the menu systems of a variety of desktop environments.

The new software update feature queries a list of mirrors for updates to packages currently installed on your system. Just select the packages you want. Click on 'update' and you're away. Well, that's the theory anyway. In practice you find that if you select a few packages, they'll be interdependant on the other ones that you didn't select (hardly Mandrake's fault) and if you select the whole lot, you'll find that one or two of them won't be found on the mirror. Or you'll discover there will be some problem downloading. Typically, these will be the packages that something else depends on...

King of the desktop

All in all though, Mandrake is as easy to use as ever, with more features than previous versions and some pretty cool custom tools. The good interaction between the Mandrake developers and the development community in general is manifestly obvious in the way various software elements are integrated so well into the Mandrake distribution.

This isn't to say that Mandrake is perfect, but it certainly still receives our endorsement for anyone wanting to run Linux on the desktop. Top stuff indeed. **LXF**

LINUX FORMAT Verdict

Ease of Use: **10/10**
 Documentation: **8/10**
 Features: **9/10**
 Value: **9/10**

Still the best of the bunch for newbies and general desktop use. Rating 9/10

A conversation with CRAIG HUNT

Nick Veitch caught up with networking supremo and Linux author, Craig Hunt, during a recent trip to London

LXF: Thanks for meeting with us. Perhaps you could tell us a little bit about your background in Linux.

CRAIG: Actually, I have a very long background in computer science. I started as a programmer and then an administrator. Then I became head of the network engineering group at the National Standards Institute. There I essentially became involved in network architecture, network design and running the network we had there day to day.

That's how I got started acquiring a deep knowledge of things like TCP/IP and knowing about networking, and it got me involved in my first book.

After that I moved into a position where I was chief of the Advanced Networks Technology division, where we did research into computer protocols

LXF: And that's where you came across Linux?

CH: Right. One of the things we found was that Linux was an excellent platform for us to use for our research projects because the source code was available to us. We could get in there and tinker and build the protocols we needed to build in the timeframes we needed to do it in. We could never have done that kind of exploration using NT where we couldn't see what was inside.

The thing that really struck me about Linux was that we had access to it, we could change what we chose and find out how things worked. Then my interest in writing naturally segued into writing about Linux.

LXF: So how long ago was it you started using Linux?

CH: I guess about four years ago. Of course I had a fifteen year background

in Unix, and I found it a very natural migration from Unix to Linux.

LXF: And you'd say that Linux was your favourite Unix-like OS?

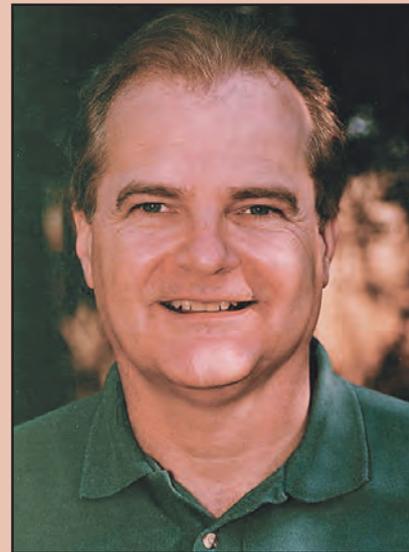
CH: Oh absolutely. Really I think that comes from the community surrounding Linux. It has more to do with the vibrancy of the Linux community; the creativity and interest. There are great systems out there like FreeBSD, but the community behind Linux is more active, and I think that's its strength.

LXF: There is a feeling about Linux where everyone has a chance to contribute.

CH: There is a structure behind Linux that comes from the GPL, and that things that are placed in it become available to everyone. It creates a sense that if you make a contribution, you are building on what other people have done, and other people will build on what you've done. It's essentially the scientific model, the way scientific research is carried out. You do something and make it openly available for peer review. You don't create something and just say it's good. Everything I've done I feel is good, it's natural. In a proprietary environment you don't have to let people look at your code and tell you whether you did an efficient or a good job. As long as it runs, and does some level of what you claimed it would in the first place, you may have a viable product.

LXF: How would you feel then if there was only one Unix-style OS, and that was Linux?

CH: Actually, I'm really opposed to the idea of 'one size fits all'. In fact, I think

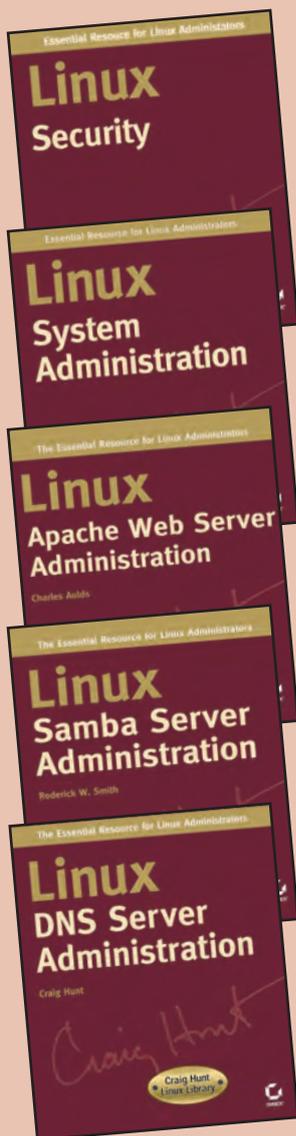


that that's one of the problems that exists with Microsoft. NT and Windows 98 aren't the same thing. Now they have converged a bit with Windows 2000, but a different skillset is still needed by the sysadmin than by the person running a desktop. Business guys somehow think they are going to save money or something if they only have one operating system

LXF: Superficially that seems to be the case.

CH: But if you were producing automobiles, and only had one source of glass for your windshields, and that factory closed down – you wouldn't be able to produce automobiles. A good business wouldn't have something in the critical supply chain come from one place. Isn't computing as important to an organisation? One virus that collapses a particular OS isn't going to effect another. The problem isn't with the technical people, it's with the business people.

LXF: Some people claim that the number of different distributions of



Linux is a handicap. Certainly in Europe, while Red Hat is still very popular, plenty of people are now using other distros.

CH: Well, in a corporate marketplace RedHat probably has a good standing because there is a company behind it, they have long standing certification processes. But there's a certain amount of the rebel in the Linux user, so I imagine the corporate popularity may make some people choose something different. I really don't think it matters that much. The underlying reliability comes from the kernel, and that's pretty much set. The name Linux is trademarked and that's a very smart move. Every bit of the code is available, but you can't use the name Linux without using an approved kernel, and that's good.

When BSD split into different fragments, they actually ended up using different kernels. There was the one using the Berkeley code, and then there was one written from scratch to avoid copyright problems. Back then AT&T had copyright, which was sold on to Novell. That caused things to go off in different directions.

LXF: Of course there are kernel forks in Linux.

CH: There are kernel forks, but they're good forks. I don't actually like to use the term forks. Linus uses the term fragments, and I like that better because fork is a technical Unix expression meaning to start a child process. I don't like using it, but I think most people might understand that term better. I went to a talk called "Good forks, Bad Forks" at the Linux exhibition back in Vegas. One guy was saying that the kernel fragmentation that's going on in embedded systems is good, and I agree.

The only bad things maybe are battles over the desktop - half the people going off with *KDE*, half going off with *GNOME*. This might dilute the perception of the desktop.

LXF: But KDE and GNOME existing is a good thing. At least it demonstrates one of the

fundamental things Linux is all about - choices.

CH: In a talk I gave the other day I said: "Linux is about choices, and that's what scares the corporate types." They don't want you to have a choice of desktop. I don't see differences in the desktop as a problem, until it comes to developing applications. Two environments isn't too many, ten may be a problem.

LXF: So which do you use?

CH: Well, I have been using *GNOME* for the book I'm currently writing, because I do a lot of Red Hat examples. For personal preference, I always liked *KDE* better. But it certainly seems to be an area of religious strife. Sometimes I wonder how good some of the applications would be if the developers spent their time coding instead of arguing. Do you ever try to stir up debates?

LXF: We don't have to, they stir themselves up.

CH: It amazes me sometimes how much people object to things. I tend to be more positive. I often look at stuff and think wow, volunteers did this. I could never have written this.

LXF: Are you enjoying working on your new series of books.

CH: Yes, I really enjoy writing about Linux, and I really wanted to write some Linux books. There are plenty of great books on DNS, but I see people recommending these books to Linux people, and you know they are going to have problems with the examples. Linux is not the same as Solaris, is not the same as HP/UX or whatever. I wanted to write a series of books that Linux users could actually read, and follow the examples, and everything would work in Linux. There isn't any point having a book that shows you how to do something on another system.

Craig Hunt's DNS Server Administration book is reviewed p30



PSST!

- Land your dream job...
- Blag what you want...
- Get away with murder...

101 INTERNET SECRETS FOR A BETTER LIFE



Get four issues free when you subscribe
Call ☎ 01458 271129 to find out more

Available with two CD-ROMs or one DVD
On sale now!

OMNIS STUDIO 3.0

You can't keep a good RAD down, apparently. **Richard Drummond** discovers what's new in the latest version of Omnis Studio.

Publisher: Omnis Software **Tel:** 01728 603011 **Web:** www.omnis-software.com **Price:** £90+VAT (Standard Edition)

Just five issues ago, we reviewed the first Linux release of *Omnis Studio*, an enterprise level, Rapid Application Development (RAD) system that enables the creation of cross-platform database-driven applications – and very impressed with it we were, too. Not content to rest on their laurels, however, Omnis Software have now come up with version 3.0. The focus for this release is serving up *Omnis* applications via the web: the server has undergone a performance overhaul and there are loads of new features in the web client. Thankfully, the fabulous Integrated Development

Environment (IDE) remains as easy-to-use as ever.

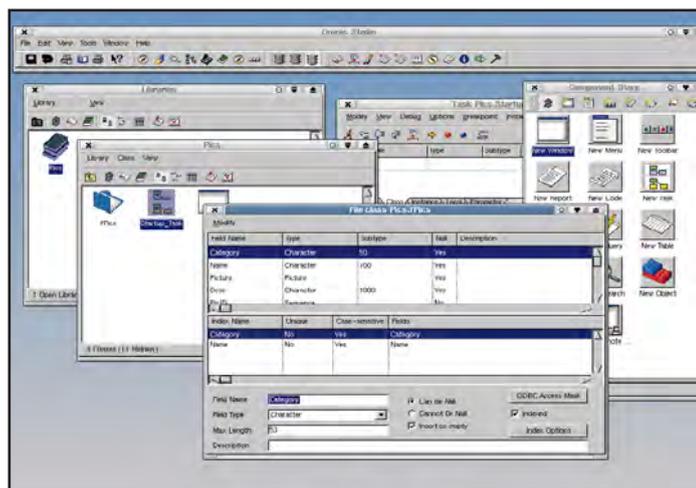
Omnis overview

Before we get to discussing the new features of version 3.0, I'll give a brief recap on what the application is all about – for the readers who missed the previous review.

In essence, the suite allows you to develop form-based interfaces for database applications. Thanks to its DAMs, or Database Access Modules,

The new version is enhanced for serving up applications on the web

it is compatible with a wide-variety of database servers including Oracle,



The drag and drop interface makes constructing an application child's play, but it would be nice to have a way of packaging them up.

Sybase and any ODBC-compliant server such as MySQL. Applications can be stand-alone, can access a remote Omnis server or be deployed through a web browser. The latter

requires a special Omnis plug-in at the client end – but only *Netscape* and *Internet Explorer* are supported – and another at the server. The server plug-in will work with *Apache*.

Omnis Studio, then, consists of four components. The Studio or development environment itself, the runtime engine and the server and client plug-in – everything you need to build and test applications. The Standard Edition is aimed at the single developer and has the single-threaded database modules from version 2.4. The Enterprise Edition has features for group development including the new multi-threaded DAMs and a concurrent version control system.

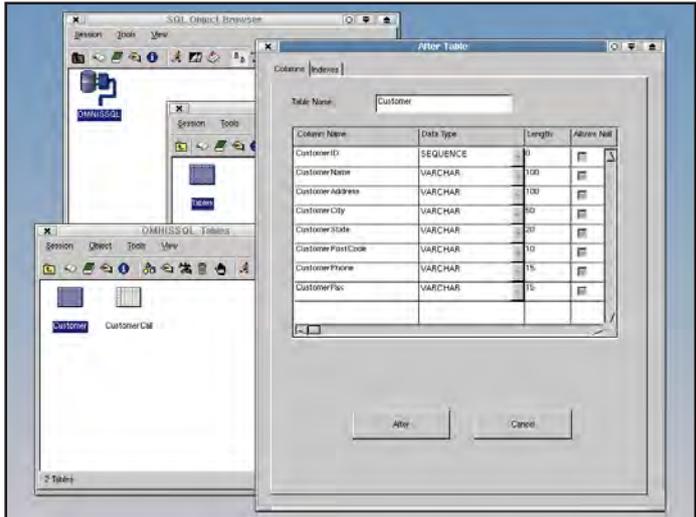
The *Omnis Studio* RAD interface is simply excellent and applications can be built with the minimum of knowledge and typing, although to do anything serious you'll need to take the time to learn the program's

object-oriented, fourth-generation scripting language. Building an application can be as simple as piecing together components with the drag-and-drop GUI and modifying their properties visually. As well as the large range of components provided, more are available from the Omnis web site and it also support ActiveX components under Windows and JavaBeans. To make things even easier, powerful wizards exist for many of the components supplied, which provide dialog-based configuration. Once the components of a form have been chosen, the overall form layout – whether for a stand-alone or web-based application – can be tweaked visually with the Forms Editor. Other great tools such as the Methods Editor mean that very little typing is required, even when you need to do your own scripting.

Omnis is provided with no printed documentation – these are available separately from Omnis Software – but it does have some excellent manuals and tutorials provided as PDFs on the CD. With these, the AppBuilder Wizard, and the online help provided, it is incredibly easy for



Omnis Software have worked hard to make the program very accessible to first time users without sacrificing power.



Omnis features a very handy integrated SQL browser which is excellent for modifying databases.

the complete novice to get their first application up and running. You'll really need to know your stuff to be able to use *Omnis* to its potential.

Web delivery

This new version is enhanced for serving up applications via the web and new features are packed into both the server and client plug-ins. The server is now multi-threaded, meaning that it can serve multiple clients more smoothly at once. The browser plug-in now has a server-side scripting capability, so that code can be executed at the client end without loading the server. This works in a similar way to JavaScript and can be used for tasks like simple calculations and interface updates.

The client plug-in allows you to display form-based applications that you create in Studio in a browser and provides a much richer interface environment than the GUI elements that standard HTML forms provide. The number of components that are usable in web-based forms has been boosted to 30 with this release and now include things like list trees, progress bars and animated GIFs.

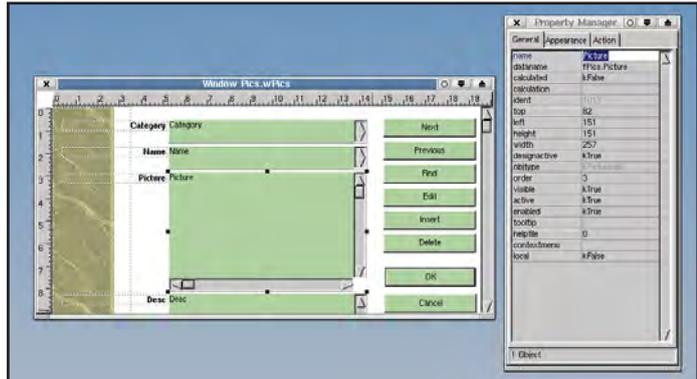
A neat feature of the new client plug-in is that it will automatically update itself and any of the components it uses if new versions are available from the Omnis website. Customers no longer have to worry about making sure they

TRY IT FOR SIZE

We have provided *Omnis Studio* on this month's coverdisc for you to try out. Turn to the CD pages on page 90 for details on how to obtain a special registration code for Omnis that will entitle you to £15 off the normal price of the Standard Edition.

have the correct plug-in version to be able to access a remote Omnis application. It's taken care of for you.

This all works quite well in practice, but I wonder why Omnis is shifting the focus to the web-serving of applications in this way. The point of being able to deliver apps via a web browser, is that it provides a wide deployment platform. This is not the case with Omnis and it's thin-client plug-in, which is available only for *IE* and *Netscape*. On Linux, the plug-in won't work on other browsers such as *Mozilla* and it won't even work with *Konqueror* which is able to use most *Netscape* plug-ins. It does make it possible for Omnis customers to embed applications in their web sites for e-commerce, but it's really only practical for intranet use – in which case, why bother with a web interface at all? *Omnis* is up against technologies like *PHP* and *Roxen* in this area as well. While neither of these have a flash GUI builder, they both make it very easy to build database applications with



Once you have created a form, you can modify its layout visually using the Forms Editor.



An example web-based application running in Netscape.

no special browser requirements, and both are open source. Java Server Pages also offer some serious competition, too. I think Omnis might want to re-evaluate their strategy here, and at the very least open-source the client plug-in.

Summing up

The improvements made in this version of *Omnis* don't seem to address any of the problems we found with the previous release and the *Omnis* suite is still not quite at home on Linux as it is on Windows. Installation still needs perfecting, for example, since *Omnis* requires write access to whichever directory it's installed to. This means if it's installed as root, you'll either have to modify some permissions or not be able to run *Omnis* under a user account. Also still a problem is the fact that there is no easy way to package up applications. It would be nice to be able to zip up all the components and data in an

application, so that they can be transported easily. At the moment, you have to install the runtime engine and deploy any library or data files required by an application separately. Lastly, the interface has glitches remaining such as windows that it will not allow you to move.

These are all fairly minor gripes, though. *Omnis* is an incredibly flexible and yet easy-to-use package that is ideal for the speedy in-house development of enterprise applications. The new web-enabled features will only serve to broaden its appeal further. **LXF**

LINUX FORMAT Verdict

Installation: **7/10**
Ease of Use: **7/10**
Documentation: **9/10**
Features: **9/10**

A new focus and more features, but Omnis remains easy-to-use and flexible.

Rating 8/10

TIME NAVIGATOR 3.5

Doing his best karaoke impression of Cher, **Richard Drummond** examines software that allows you to turn back time – but not literally!

Publisher: Quadratec Software **Web:** www.quadratec-software.com **Tel:** 01344 307420 **Price:** See boxout

The amount of data that companies have to deal with is rising exponentially. To protect this data, it is essential to have a back up solution that is designed to cope. One such solution is *Time Navigator*.

This is a cross-platform, client-server backup system. It supports both full and incremental backups and archiving, a huge range of tape and disk backup media, and a wide-variety of operating systems, including most Unices, Windows 95/98/NT, VMS and OS/2. *Time Navigator* can also backup data from certain applications directly – such as Oracle and Sybase – and understands NDMP (the Network Data Management Protocol). Both NAS (Network Attached Storage) and SAN (Storage Area Network) architectures are supported, and data can be backed up either in a proprietary or tar and cpio formats.

Quadratec claim that ease-of-use is a major selling point of *Time Navigator*, and the Unix and Linux versions are configured and controlled with a Tcl/Tk-based interface. The 'Time' part of the title stems from the fact that users can restore their system from a backup to the state it was at some arbitrary time in the past.

Backup time

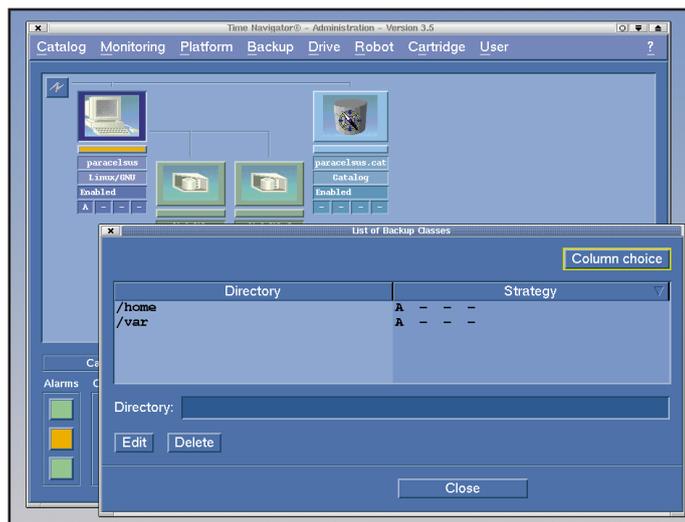
Installing *Time Navigator* is fairly straightforward. Your network must have at least one machine with the *Navigator* server installed on it, and each machine with data has to have

the client software installed. A server is typically a machine with backup devices attached, but you can configure 'robots', that is, remote backup devices which are controlled from a server. You can also share pools of devices between servers.

Installing a *Time Navigator* server or client is fairly straightforward. First create a user account called 'tina' to act as a super-user account for the application. Then execute the install script, choosing whether you want the client or server and entering the install wizard.

Once everything's installed, it's time to set up a policy for your network. *Time Navigator* defines the various components and back up tasks within your network as objects. So, there are hosts (machines which have data you want to store), classes (specific directory tree with data you want to back up), strategies (a policy of what to back up and when), drives, robots and so on.

The Admin GUI provides a pictorial view of these elements within your network and allows the administrator to set everything up fairly quickly – providing they know what they're doing. Adding a host to the network is simple, since the software automatically searches for machines with a client installed. Then, each host can have up to four backup strategies, each defining a set up classes to back up to a pool of storage devices. A strategy covers both full and incremental policies. You can filter which files in a directory



Before you start your backup, you need to define what devices are available and what you want to save.

tree will be included in a backup class by regular expression, file size and modification date. As your network grows, it is important to grow your backup strategy and *Time Navigator* allows for this.

The users on your network operate *Navigator* through a lesser GUI. This does allow them to add their own classes to the backup strategies for their host, but little other configuration. The main focus of this interface is recovery. It provides powerful filtering and viewing controls for users to browse through their filesystem as it was in any point in time covered by a backup. The user can step forwards and backwards in time and filter the files shown by regular expression, modification date and size – and choose to restore any selection of these to their system.

Time Navigator provides some powerful services to automate and manage the backup and archiving of a network of machines. Even though it is aimed at the enterprise level, I cannot help but feel that its interface

PRICE

Pricing for *Time Navigator* depends on the configuration required. The basic cost is £3200 for the server and £320 per client. You'll have to add to that the cost of libraries to drive the types of backup devices you wish to use. Special deals are available. A license for one Linux server, 10 clients and an 8 slot library is £3200. Contact Quadratec for further details.

is a lot more complex than it needs to be. The admin and user GUIs are not particularly intuitive and provide little in the way of on-line help. Documentation is, however, provided on the CD in PDF form and this explains the concepts and operation of the package, but isn't readily digestible by the beginner. Still, if your business has someone willing and able to learn how to use *Time Navigator*, it will provide a reliable and flexible back-up solution. **LXF**

LINUX FORMAT Verdict

Features: **9/10**
Ease of use: **6/10**
Documentation: **7/10**

A serious backup system for enterprise applications, but not very user-friendly.

Rating 7/10

OPERA 4.0 Beta 3

Will the Linux port of this web browser be good enough to justify actually paying for it? **Richard Drummond** investigates.

Publisher: Opera Software **Web:** www.opera.com

With all the media attention that has been focused on *Mozilla* and *Konqueror*, you could be forgiven for ignoring what's been going on at Opera Software in Oslo, Norway. There they have been quietly beaver away on the Linux port of their award-winning browser of the same name, releasing test versions every three to four weeks. With the recent third beta version, *Opera* is now nearly feature complete, so we thought we'd take a look and see how it stacks up against the competition and whether, as the only Linux browser with a price tag, it has anything to offer.

What is Opera?

Opera Software have been developing *Opera* on Windows for five years, and although it has a pitiful market share when compared to *Internet Explorer* and *Netscape*, it does have devoted users who are quiet willing to pay for

a browser that is light, fast and stable – not standard features of either of the big two freely-available browsers. In 1998, hoping that users on others platforms would also be willing to stump up the cash for a better browsing experience, *Opera* started a project to port their browser to other operating systems including Linux, MacOS and BeOS.

The newest Windows release of *Opera* is version 5.0, but version 4.0 still packs a very full set of features, including some not found on rival

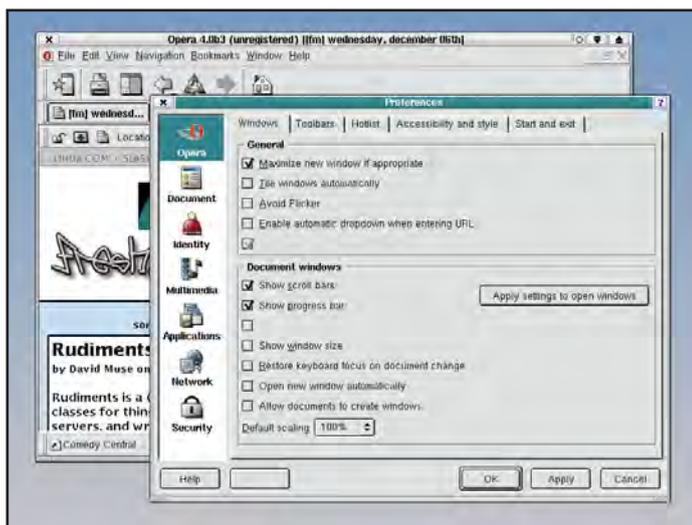
In 1998, hoping users would be willing to stump up cash for a better browser, Opera started to port their browser to other OSs

browsers. At the top of the list is standards compliance with support for HTML 3.2 and 4.0, XML and WML (for browsing WAP sites), HTTP 1.0 and 1.1 (including resuming of

GETTING OPERA

You can download the latest beta version of *Opera* for free via the Opera website at www.opera.com. At the moment, the beta versions are time-limited to 30 days use and, once the evaluation period is over, you will either have to download a new preview version or pay a registration fee if you continue to use it. What the charging policy will be when the finished article is released remains to be seen.

Opera is available in several versions for x86 and PPC architectures and as RPM and Debian packages. It's built with version 2.2.1 of the Qt toolkit, and you can download packages which are either dynamically or statically linked with the Qt libraries. If you have the relevant version of Qt installed on your system (it's used in KDE 2.0), then the former is the one to go for since it's a smaller download.



Virtually every element of the browser can be configured.

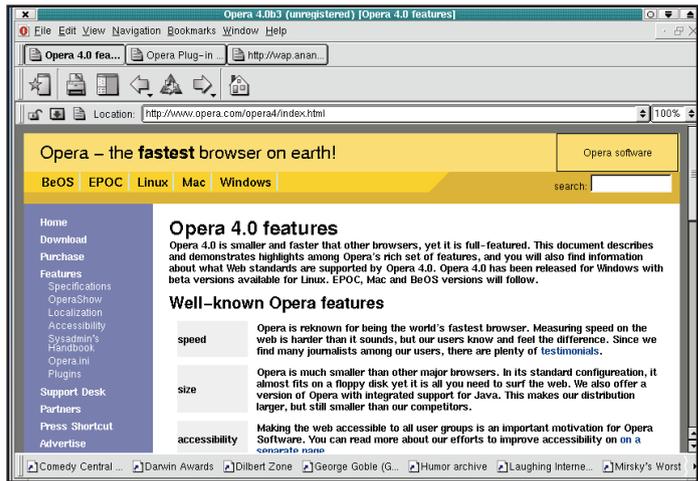
downloads), SSL, Javascript (fully ECMA-compliant) and Java. Other notable items include full support for cookies, an integrated POP3 mail client, and a Netscape-compatible plug-in system. There are also neat features that I have previously only seen in Amiga browsers – and Windows curiosity *Neoplanet* – like a transfer window where the status of all files currently being downloaded can be viewed and modified and an integrated cache browser which lets you pull up previously viewed pages from your hard disk.

The Linux port of *Opera 4.0* is well into the beta testing, with the third beta release providing most of the core browser features. In fact, users who have tried any of the earlier test versions will be pleasantly surprised at just how complete it is. There is no print facility yet, no Java, no plug-ins, no documentation and some of the configuration options aren't complete, but pretty much everything else is functional.

Visiting the Opera

When you boot up *Opera* in Linux, the first thing you'll notice is how quickly it starts. The second thing you'll notice is that its interface looks a bit odd when compared to other browsers. Don't worry about the layout, however, as you can configure just about everything to look the way you want it to. All the major controls are on tear-off strips that you can move around or hide as you please and some components, like the hotlist window, can be docked or left floating. *Opera* isn't as theme-able as *Mozilla*, but it is more easily customised. The preferences will let you choose which button images are used in the navigation bar and which fonts, colours and Qt theme will be used in widgets. It's not just the interface that can be configured, however, but just about every other feature of *Opera* as well, including every last detail of how HTML pages are rendered.

One thing that takes a bit of getting used to is the way *Opera* supports browsing multiple pages. Rather than displaying these in separate X windows, there is only one



Opera aren't reticent when it comes to blowing their own trumpet.

main *Opera* window and each page gets its own embedded window within it. You can move and size windows around and expand a page to fill the whole window, just like you would under Windows, plus there are controls to tile and cascade them. Initially I found this behaviour rather disconcerting, because I loathe this type of interface in Windows, but after persevering with it, it turned out to be less of a distraction than I had expected. I would still much prefer the option of being able to display separate pages in separate physical windows, though.

Minor differences like these aside, *Opera* is used identically to any other browser and should be instantly familiar to most users. You just get a good deal more control in here than you do in most other browsers.

Opera in action

So, how does it perform? Well, *Opera* boldly claim that this is the world's fastest browser, and it is certainly no slouch at rendering HTML pages, even complex ones with nested tables. On a machine with a fast Internet connection it can download and layout Freshmeat, for example, in under four seconds. For comparison, *Konqueror* manages about the same,



One of Mr Drummond's favourite haunts shows off the browser's 'odd' window management system.

but with *Netscape* and *Mozilla*, you can go away and make a cup of coffee and they'll still be grinding away when you come back.

What really makes *Opera* stand out from the competition, however, is the speed of the user interface. New windows open immediately – not with a two-second lag like they do in *Konqueror* – and there's none of that silly busy-pointer nonsense. *Opera* is always waiting for your command. This comes into its own when you have limited bandwidth such as a modem connection. Pages in *Opera* always seem to be usable more quickly. This is partly because *Opera* does a better job of rendering

THE BIG QUESTION

You might think that *Opera* Software is mad to actually try to sell a web browser, especially in a market like Linux where there is a reluctance to pay for any software. And no doubt many of the Linux old guard will dismiss *Opera* entirely simply because it is proprietary software; they're entitled to their opinions. The question is, though, how important to you is the quality of your browser? Personally, I use a web browser enough that the unique features

that *Opera* provides are worth shelling out the money. Of course, I would prefer an open source solution, but there are times when pragmatism outweighs idealism. Short of starting my own browser project, I don't see the open-source arena coming up with a browser that fits my needs as closely as *Opera* does. Hopefully, enough Linux users will think this way to encourage *Opera* Software to continue development of this excellent product.

Konqueror. Of course, the lack of Java and any plug-ins means there are a number of sites you just won't be able to visit yet.

Opera is not without its faults; I have come across several web sites that it refuses to cooperate with. As an example, with Talk21, BT's webmail service, it won't allow a user to log in, and Abbey National's e-banking service causes it to throw segmentation faults. Admittedly, both these sites are unusable in *Konqueror* and unreliable in *Mozilla*, but *Netscape* copes in its usual klunky way. There are many other more obvious bugs in the current beta of *Opera* – for example, if you interrupt a transfer you sometimes need to restart the program before the browser will connect to another page – and clearly there is room for improvement with regard to stability, but these issues will be addressed as the beta program continues. *Opera* on Windows is rock solid, and there is no good reason why it shouldn't be as equally bomb-proof on Linux, too.

At the moment, *Opera* for Linux is an interesting preview rather than a useful browser. But once the bugs are squashed, it will offer comfortable and efficient browsing and a degree of control that is unequalled by any current Linux browser. All I can say is, roll on the next beta! **LXF**

LINUX FORMAT Verdict

Fast, compliant and reliable, we can't wait to get hold of the finished product.

SHOGO Mobile Armor Division

Who will come up trumps in the battle of the mech games? **Paul Cavanagh** road tests a metal monster in the first of this month's monumental megamech reviews.

Publisher: Hyperion **Web:** www.hyperion-software.com

Price: £29.99

The future will see military hard-cases and research and development boffins sniggering into their hands at the very thought of a tank. Using tracks to get about? You must be kidding. What you want is a great huge mechanoid that runs about on very human-esque legs, blasting everything in sight with a variety of ordinance. How do I know this? Well, my friends, I have seen the future... and it's robot shaped!

Unfortunate Timing

Alright, that's a complete load of cobblers, I haven't seen the future at all, I've just been playing some mech games and got into it all a bit too much. Mind you, it's hard not too, and that's a good thing. Less good, at least as far as Hyperion are concerned, is that fortune and circumstance dictated that two games featuring mechs have come out at the same time. Given that both games are at heart first person shooters, and both give you a large robot that can turn into a vehicle and



This is where you select your mech. You flip a switch and a computer voice tells you all about your choice.

whose primary function is to shoot things, it would be churlish not to compare the two games. *Shogo* is a very good game indeed, but when placed next to the superb *Heavy Gear II*, it doesn't seem quite so impressive. For the moment though, let's just look at what *Shogo* is all about, and find out its strengths before making any direct comparisons.

At the beginning of the game, you start off in your barracks without your mech, and you can wander about the space-station and admire the view out of the windows. There are lots of people here, drinking in the bar or hard at work using computers or doing maintenance jobs. Although none of them will fire on you, none of them seem to be particularly friendly. The only thing they're likely to say to you is "Watch it pal!" if you bump into them, apart from the barmaid who tells you that she can't serve you while on duty. Still, this is a

nice gentle way to ease you into the game, and get you used to the controls. There's an information network you can



Brrr! Why can't they put a decent heater in these mechs? Good job I brought my Fisherman's Friends with me.

use to pick up on some background facts (not that you'd need to really, it's the same old story – terrorist threat

you're transported to a planet's surface where you can begin to familiarise yourself with your

Terrorist threat blah blah, big chief is a psycho blah blah, shoot, shoot and shoot again

blah blah, big chief is a psycho blah blah, shoot, shoot and shoot again etc.) When you've had enough of milling about aimlessly you can meet your commanding officer and receive the first mission briefing.

Exmechs Shopping

Now you can choose your mech. There's only a few models to choose from, and you can opt for models that excel in speed, armour, weaponry or general all-round abilities. Once you've chosen your mech, you're stuck with it for the whole game, so be careful. Next,

machine's abilities. Basically, it's very much like controlling a human in a first person shooter – you can duck, crawl and jump. When you

look in a certain direction with the mouse, the mech's whole body turns in the same direction. Your mech can convert (Transformers! Robots in disguise!) into a high speed armoured vehicle, but while in this state, you can't shoot. I've only used this function on a couple of occasions, but I guess it might be handy for retreat – I'm no lily-livered coward y'know, I stand my ground and fight to the death (usually mine).

Early on in the game, you will only encounter artillery and light armoured vehicles – these are relatively easy to destroy, but difficult

**Snipe him!
Zoom in
and
watch
him
crumple
when you
pull the
trigger.**



to target because of their small size. There is a zoom facility on some of the weapons which helps a lot. Later on you'll encounter other mechs, which are a complete pain – the best bet is to hide somewhere and take pot-shots. oh, and if you come across any civilian vehicles, it's a good idea to blow them up because they often contain armour or energy pickups.

Thigh-high Skyscrapers

The game environments are interesting and detailed. First off, you'll be in a city, and the perspective is distinctive because the buildings only come up to waist-height. Later, you'll encounter outdoor arenas, which can prove challenging when facing multiple enemies – again the best option here is to use the sniper-rifle and pick off the bad guys from a distance. Sooner or later, though, you'll be inside a building complex where you'll encounter the more usual tunnel systems that I've come to expect in first person shooters. On the whole though, the level design is fairly interesting and challenging.

Nice Briefs

Each level has a number of objectives, and your mission brief is revised as you go along. You are in constant contact with base, and they'll update you as necessary. This is a good way to structure the game because you don't know what to expect when you enter a level, and you do feel as though you are on a real mission. There are a number of contacts at the base, all of whom have different characteristics. Your character has a different relationship with each of these people and will react accordingly to their

communications. This does add a little bit of narrative to the game but, to be honest, the plot is very much secondary to the action.

Gosh, it's hard!

As you progress, it becomes apparent what a challenge the game is. Bad guys frequently ambush you, hiding just around blind corners, or at the top of ladders or through doors. Sometimes they'll sneak up behind you. Given this, you'll soon be biting your nails, and rushing through doors guns blazing, hoping to get that vital first shot in before you're caught in a hail of gun-fire. After a few levels you have to abandon the mech and progress on foot, and it's amazing how vulnerable this makes you feel. Where your mech can have armour and energy ratings in the thousands, without it, you've got a maximum energy rating of a hundred, and that can go down to nearly zero after one good hit. Nasty. *Shogo's* difficulty (at least as far as I'm concerned) is one of its major strengths. It's good to be on edge in a game, to feel the fear, to resort to saving after every encounter with the enemy. Sometimes things can become very desperate indeed, as you try and plug someone when you've only got ten strength units and no armour left. If you've got the guts, you can get through it all.

Graphically, the game is well up to par with its contemporaries, everything is remarkably well detailed and the weapons, enemies and explosions all look stunning. There are three graphical detail settings, and for the most part I was using the lowest available, which speeds things up some, but even then the game looks

PUT DOWN YOUR WEAPONS... PLEASE!

OK, so it's the most popular and most common style of game available on the PC, but do we really need any more first person shooters on Linux? Well, yes, we're looking forward to *Deus Ex* (which is sooo much more than an FPS), and it would be nice to see *Half Life*, but there's more to life, you know. What about some stealth games? *Thief* is a great variant on the FPS where the object is to avoid getting into fights, and *Metal Gear Solid* is arguably one of the best games ever released. What about some 3D racing games? Not really as popular on the PC as the PlayStation, it's highly unlikely that we'll see versions of the *Ridge Racer* series, or

Gran Turismo, but surely we could have something to sate our need for speed. How about a 3D platformer? Something to rival *Sonic* or *Spiro* would be great. Variety, after all, is the spice of life.



Deus Ex: another FPS, but with oh so much style.



Big meaty weapons are always the most fun. Who mentioned Freud?

lovely. The music and sound effects are also up to scratch, and the clear speech leaves you in no doubt as to what you're supposed to be doing at any given time.

The selection of weapons is pretty much bog-standard for first person shooters – there's nothing here you wouldn't find elsewhere, which is a bit of a shame. One oversight is that all the weapons are hand-held, which seems a bit daft when you've got a great big mech to play with – why not put rocket launchers on the shoulders? This is just one of the ways in which *Metal Gear II* scores points over *Shogo*.

All in all, this is a splendid first person shooter, maybe not as good as *Soldier of Fortune*, but it's up there with the best. The fact that you're fragile body is surrounded by a

thousand tons of metal rather than a khaki uniform and a bit of Kevlar doesn't make that much difference to gameplay. True, you can change into a car, and the perspective is different, but it really doesn't change the fact that this is plainly just another first person shooter. So take your pick from the growing pile of such games available on Linux – this one is neither much better or much worse than all the others out there. **LXF**

LINUX FORMAT Verdict

Features: 7/10
Gameplay: 8/10
Value: 7/10

Competent, if a little uninspiring, first person shooter.

Rating 8/10

HEAVY GEAR II

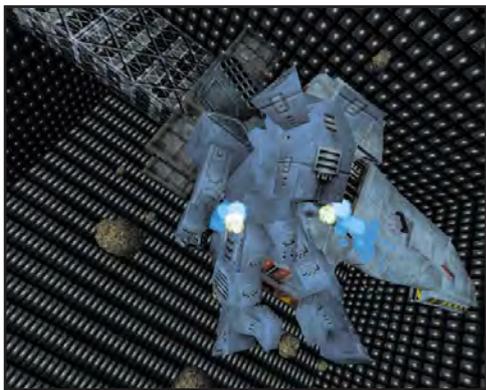
What's this then? Oh, it's just another mech first person shooter. **Paul Cavanagh** discovers just how wrong first impressions can be...

Publisher: Loki **Web:** www.lokigames.com **Price:** £ 39.99

On the face of it, this could easily be *Shogo* relabelled – giant robots blasting everything in sight, a plotline centring on an age-old war, primarily a first person shooter. Good fun, but nothing really special. But load it up, have a fiddle around and you'll find that this game has hidden depths going far beyond the standard shooting experience. One thing that's nearly always going to impress me is a hefty manual – it's a hint that a game has more than meets the eye. While it is possible to pick up and play *Heavy Gear II*, a quick browse through the 112 page manual will definitely help ease you into the action. A good deal of the manual is waffle, dealing with pilot biographies and the various mechs available, but the rest of it is a very useful guide to the gameplay.

All things to all men

Before you get put off, don't worry, this isn't the sort of game where



control system. On the whole, everything is fairly normal here, you can walk, run, jump, strafe, kneel and lie down on the floor (OK the last one isn't that common,

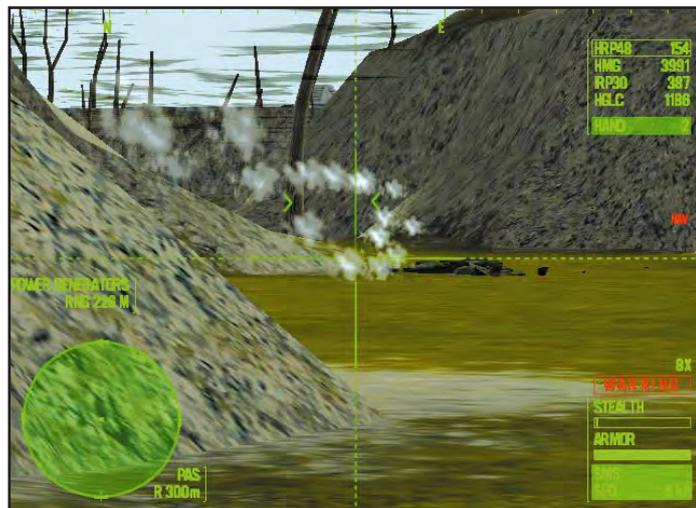
This is a training mission in a zero-g simulator. You can see the gear's retro rockets in action.

you'll have to spend hours fiddling about to get going – this is without a doubt an all action blaster, and a damn good one at that.

Those of you who don't fancy getting involved in strategic gameplay can go right ahead and blast away and you'll have fun doing it – you just need to spend a little time getting used to the controls. But

those gamers who prefer a more cerebral experience will enjoy the finer points of the game – and will be all the richer in terms of enjoyment and accomplishment for it. But I'm getting ahead of myself here – I'm just trying to point out that this is a very rare thing in the gaming world – something that would appeal to just about anybody.

It really isn't easy to know where to begin in describing what sets this game apart from the rest, but you've got to begin somewhere, so I'll start with the



Using the zoom feature allows you to pick off enemies before they even see you. Note the smoke, it's much better when it's moving.

but you get my drift). What makes this game that bit more interesting is that, just like a person, your mech can pivot at the waist so you can, say, look to the left while you're walking forward. This might seem like a minor

This is a very rare thing – something that would appeal to just about anybody

detail, but it really does make a huge difference in the way the game feels.

In a battle in can be very useful to be running towards a goal and being able to swing your aim from left to right, picking off enemies as you go. It can get confusing because you can forget which direction your feet are pointing in, but your heads up display gives you a hand by providing a set of brackets. When your crosshair is centred in between the brackets, you're facing the same direction as you are travelling in. Ace. All of the movement controls can be found on the number pad of your keyboard, while the direction you face

is controlled with the mouse. It takes a little getting used to, but once you've sussed it, you'll wonder how you ever coped with just those cursor keys. In addition to this, your mech can travel on roller skates. Have you seen those trainers with the little wheels that flip out? I reckon the inventor got his idea from this game, because that's how it works.

If you've got some travelling to do on relatively flat ground, this is the ideal way to do it. It really is fantastic, just rolling along watching the beautiful world go by.

Plan ahead

Unlike most FPSs, you can't pick up ammo, weapons and health as you go along. At the start of each mission you get the opportunity of modifying your gear (that's your big robot). At this stage you're given a briefing for the mission and can decide what weapons would serve you best. You can take your pick from over 70 weapons grouped into four main

categories – projectile hunks of metal (bullets) that you can shoot; Missiles which are self-propelled; indirect weapons such as mortars and grenade launchers; and energy weapons which are either laser guns or hand to hand vibrating weapons like axes or swords. In addition you can alter your gear's shields, mobility and stealth settings. This is a great way of organising things, because it allows you to experiment with all of the game's weapons from the outset, and once you've played a mission, you'll have a better idea of what you'll need to take with you in order to survive. You won't be able to take all the weapons with you on any missions, you're limited by what you can carry and by your mech's threat value – the better the weapons the higher your threat value. You are assigned a maximum threat value for each mission. Equipping a mech is fun and can be a real juggling act when you've got a small threat value to work with.

Once you've started a campaign game, it's very difficult to stop playing. As you look about you you'll be amazed by the crisp clarity of the visuals around you – the levels vary wildly in appearance, from city slums to arctic forests, from verdant pasture-land to the vacuum of space. Oh yes, space. In space your mech is powered by little rockets, and you can do the whole travel through 360 degrees bit. Very challenging indeed, but it's a hell of a lot of fun blasting the enemy when you're upside down. It certainly adds a twist to the already novel gameplay.

Ooooooh, pretty

The graphics throughout are quite



Some levels will take of a lot of practice before you crack them. As always it helps if you know where your enemy will be coming from –

Here's the tactical map, showing your squad and some of the foolhardy orders you can issue.

simply staggering, some of the levels are real works of art. This game certainly does not suffer from the monotony of underground tunnel systems that have begun to be the bane of my life, and it's such a relief to be able to engage in meaningful battle under a blue sky surrounded by mountains and lakes. This visual excellence applies to everything in the game – structures, enemies, smoke, fire, explosions, the sky, the stars, your own mech, everything looks absolutely gorgeous.



Shoot! shoot! shoot! shoot! shoot! shoot... Oh, it's a tree. Shoot it anyway, it could have a gun! Shoot! shoot! shoot! shoot! etc.

Each mission comes with a briefing in the form of a short animation which you can watch before embarkation, and as you achieve your objectives you will be updated by your commanders. As you engage the enemy you'll notice that they react intelligently to you. They'll try to run when injured, but when healthy they are damned aggressive. There are all sorts of enemies to encounter, mostly infantry, tanks and other mechs, and just like everything else they look great and act fast.

especially because your mech has got a telescopic sight, so if you know where they're hiding you can hit them before they know you're even there.

After you've played a few missions, you'll be introduced to another very important element of the game – your squad. Once you've proved yourself, you'll be able to take charge of up to four other mechs. As a squad leader, you can issue commands to the other mechs in your team (don't call them cannon fodder!). This is very simple to do, you

can either go to the superb overhead map where you can tell your squad or any member of it to move to a certain location by clicking on the map. You can also dictate which formation they move in, tell them to attack or defend a target or to cease fire. You could – if that was your 'thing' – play the entire game from the map screen (you can move your own mech from here too), and treat the whole thing as a strategy game. Sometimes, you'll want to issue an order quickly, without bringing up the map and you can do this by using simple keyboard shortcuts (ctrl+r is regroup, ctrl+t is attack my target etc). How much you get involved in these tactics is really up to you. Left to their own devices your squad will do the sensible thing and follow you about, attacking any enemies they come across, but they'd do a lot better under your guidance. For example: you're up against four tanks – three little ones and one big one. It

makes a lot more sense to get your squad-mates to all attack the biggest strongest vessel together than randomly attacking all four, you'll save your energy and ammo. An order is simple to issue, but beware, some of your mech pilots are a bit dodgy, and if they don't trust you they probably won't obey your orders. Strategists will love this element of the game, while wannabe action-heroes can forget it even exists. Perfect.

Steady now...

Frankly, I'm never going to have enough space here to describe what makes this game so damn good. I've only been talking about the mission-based campaign games here, but there's plenty of other gaming styles to keep you happy. There's training missions to teach you everything you need to know, there's historical missions where you can choose to take part in a celebrated battle in the long running war, there's instant action where you can choose up to four squad-mates and four enemies to battle out with over various scenarios and there's a multiplayer version which I just haven't had time to play. Value for money or what?

So here's the crunch. Games that deserve a perfect ten really are as rare as Brigadoon. Given the fact that I can't think of a single bad thing to say about *Heavy Gear II*, and that I'm determined to play it to the bitter end no matter what, means I'd be happy to say that this, my friends, is a perfect game. Without a doubt the very best game available for Linux users at the moment. And you get a fantastic postcard depicting Tux dragging Bill Gates up the Empire State Building to boot. What more could you possibly want? **LXF**

LINUX FORMAT Verdict

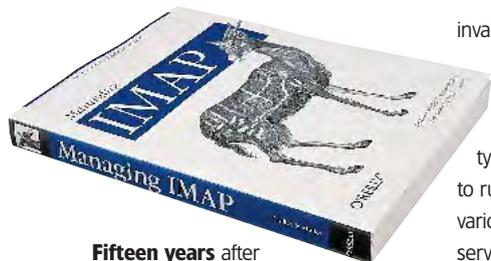
Features: 10/10
Gameplay: 10/10
Value: 10/10

Check it out! Ten out of ten all round! Guess what the rating's going to be. Don't be daft, buy this game.

Rating 10/10

Managing IMAP USING THE INTERNET

Publisher: O'Reilly **ISBN:** 0-596-00012-X
Author: Dianna Mullet and Kevin Mullet **Price:** £21.95



Fifteen years after its inception, the Internet Mail Access Protocol (IMAP) is at last beginning to displace the simpler Post Office Protocol (POP) as the preferred system for email hosts. *Managing IMAP* gives a complete guide to this email revolution for everybody from the user to the administrator whatever the OS.

The book is divided in six parts. The first section gives an overview of email generally and IMAP in particular; discussing the different ways IMAP may be used and comparing it to POP3. Part two is an overview of some of the more popular IMAP clients, listing the features you should look for in a client. The next two sections give detailed guides to installing, configuring and administering two of the most popular IMAP servers, *Cyrus* and the *UW* (*University of Washington*) server. These chapters will be

invaluable when choosing a server, since they discuss the suitability of each server for different types of installation and the typical hardware resources required to run them. Part five is a melange of various topics, including security, server-side mail filtering and tuning the server. Last but not least are the Appendices, with Appendix C giving a guide to the commands that make up the protocol itself.

This is an enlightening read; the authors clearly know their stuff and have presented a well-written and thoroughly digestible guide to the world of IMAP. Minor grumbles include the fact that the section on IMAP clients doesn't mention any non-web-based open-source clients and that a companion CD giving some of the software mentioned in the book and the protocol specs would have made a useful addition.

Richard Drummond

LINUX FORMAT Verdict

If you are planning on joining the email revolution, this is the definitive guide.

Rating 8/10

Publisher: Computer Step **ISBN:** 1-84078-110-6
Author: Brian Austin **Price:** £4.99



Windows dominates the computer world. This is a sad, but nonetheless true, fact. However, it's not THE ONLY operating system in the world, and *Internet Explorer* is not the only browser. We just had to clarify that because, reading Brian Austin's otherwise well put together book, you could be forgiven for thinking you were reading *Using Internet Explorer*. In fact, if you excised the tiny Netscape Communicator section (which mentions IE almost as much as Netscape) and the screenshot from *NetObjects Fusion* on page 347 that's exactly what you'd have. And Linux... what's Linux?

For a new Windows user, this could be the ideal introduction to the 'net: it has sections on choosing hardware, configuring the OS and making that first vital connection. But the book's MS bias really colours everything about it – even Macs are given short shrift – so the book's useful stuff like advice about keeping emails small, basic web page design

and safe surfing are buried amid all the DUN configuring and changing your (IE) homepage.

The final section features the now ubiquitous selection of world wide web links covering topics from animals and auctions to shopping, and cyberculture.

The only reason you might want this is if you're dual-booting and using Windows to surf. But if you are, you'll probably find the text quite a bit simplistic and the lists of links something you could easily find by actually firing up your browser.

Andy Channelle

LINUX FORMAT Verdict

Alright for Windows users, but this is a Linux magazine!

Rating 2/10

Internet Shopping

Publisher: Computer Step **ISBN:** 1-84078-107-6
Author: Geoff Preston **Price:** £4.99

We love shopping, and the prospect of doing everything online without the endless trudge through a crowded highstreet sounds like sheer bliss. But can a book like *Internet Shopping* take the place of portals such as Yahoo Shopping, thevirtualmall.co.uk or free2give.co.uk?

Actually, it can. In addition to the useful stuff about online security, cookies and returning goods, there is something to be said for having a selection of ordered and tested sites broken into easy to find categories. Of course, you could get all this with a search engine, but you'd probably

have to wade through an awful lot of rubbish to a) find the right product and b) get it delivered in the UK.

Internet Shopping covers sites selling everything from tableware to evening wear and also contains some useful information regarding your activities and rights online. Each entry has a brief description of what the site is about, what you can expect to find there and, occasionally, an opinion on the retailers prices.

The comprehensive listings also cover what you'll need to keep your software and peripherals up to date (though with a Windows bias,

naturally) and some of the more popular auction sites.

The biggest problem with printed material is that it takes so long to prepare and update, so where a dedicated shopping web site could be bang up to date, book readers may come across the occasional dead link.

Andy Channelle

LINUX FORMAT Verdict

Very good for online shoppers who know what they want, but don't know where to get it.

Rating 6/10

Web Application Development with PHP 4.0

Publisher: New Riders Publishing **ISBN:** 0735709971
Author: Tobias Ratschiller and Till Gerken **Price:** £29.99

With the release of PHP version 4 earlier this year, the number of books covering the subject has increased dramatically.

If you're a budding ASP developer or just want to tinker with PHP, but don't yet know the first thing about it, this certainly isn't the book for you. Any programming book which discusses code comments in the second chapter isn't for those with a faint disposition or novice users.

This book expects you to at least know how to setup PHP and Apache, as well as having a rather in-depth understanding of the PHP syntax and internal functions. If you don't know how to build MySQL support into PHP, or make it run as a binary, you'll need to look elsewhere for that kind of information.

All is not lost, though. For experienced PHP users, particularly those who want to expand on their knowledge, but don't quite know how, there are interesting, and occasionally rather innovative and obscure, ideas and code snippets which are easily adapted to any situation or code style. It is certainly more of a concept and guidance text, rather than many other tutorial styled books, as you can just jump in at the beginning of

almost any chapter and look over what is covered. Each chapter contains a selection of code and overviews of what is happening, although more often than not, you'll need a three page library printout, followed by an inspection of each function, so you end up flicking back and forth to try to remember exactly what the function contained. Fortunately, the code is also on the accompanying CD, so at a push, you could look at the listings onscreen and read the text in the book.

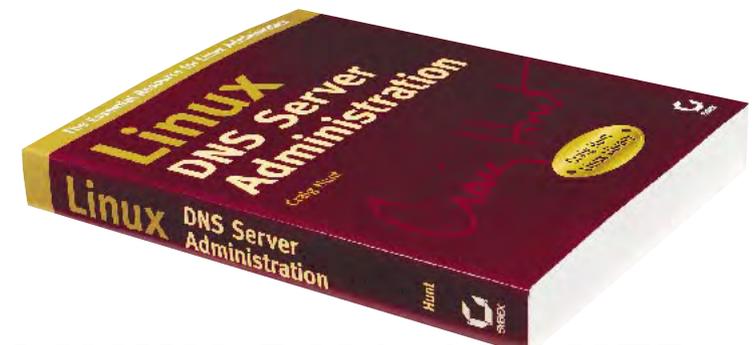
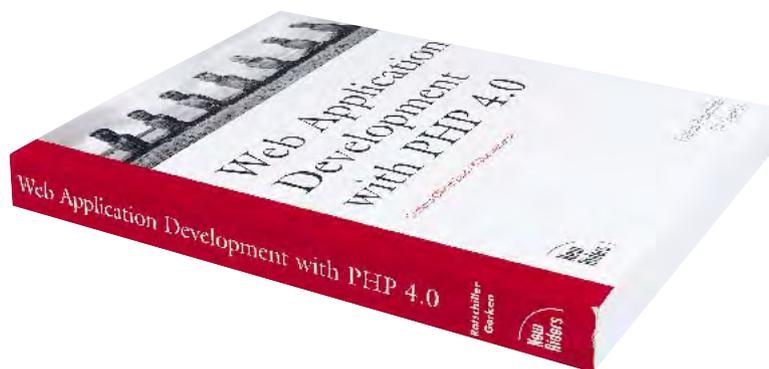
I'd personally expect a bit more for my money. This is a bit thin, and given the level of expertise it is aimed at, there isn't really a great deal which many people seriously involved with using PHP could not pick up from web sites such as phpbuilder.net or from the various projects available on the Internet.

Dave Coulson

LINUX FORMAT Verdict

Interesting book for experienced PHP users, although they could probably work most of it out for themselves in the time it takes to read the book.

Rating 7/10



LINUX DNS SERVER ADMINISTRATION

Publisher: Sybex **ISBN:** 0-7821-2736-3
Author: Craig Hunt **Price:** £29.99

The Domain Name System (DNS) is the glue that holds the Internet together. Without it, finding a particular host amongst the tens of millions out there would be impossible. Craig Hunt's book aims to be a complete reference for those that have to administer DNS servers on Linux systems.

DNS and BIND by Albitz and Liu has long been considered the bible for DNS administrators, and Hunt's book differs from that title in two ways. Firstly it concerns itself only with Linux systems; not having to include information on all the different flavours of UNIX gives the book a clearer, less cluttered feel. Secondly Hunt only deals with BIND version 8, though there is an appendix covering the forthcoming version 9. All recent Linux distros ship with BIND V8 so you should find that your system is covered.

However, there are still quite a few systems, Linux and otherwise, out there that are still running BIND V4.9.x. and there are significant differences between the two versions. In particular the configuration files are completely different. A section on upgrading to BIND V8 would have been useful.

The book is divided into four parts; How Things Work, Essential Configuration, Advanced Configuration and Maintaining a Healthy System. Many sysadmins with relatively simple systems could get away with only

reading Part 2. Here we are taken through the configuration of the resolver, a caching and slave server and a master server. As well as straight editing of the text files Hunt covers GUI tools such as tksysv and linuxconf.

The third part of the book deals with topics such as creating sub-domains, delegation, reverse domains etc. Here the description and examples show what the sysadmin needs to do, and what to avoid, in more complex situations. The tricky subject of delegating reverse domains is handled clearly and concisely.

Testing, security and logging are all covered in part 4. We are taken through various techniques to increase security; including restricting transfers and queries, encryption, authentication, integrity etc.

This is a fair introduction but may not quite be meaty enough for some administrators needs. The section on testing is more complete, containing a good description of the tools host, dig and nslookup. Power users may need something more heavyweight.

Steve Heaven **LXI**

LINUX FORMAT Verdict

A useful book for the majority of DNS administrators. If however you have hundreds of domains spread across many servers then you probably need something a bit heavier.

Rating 7/10

IRC ROUNDUP

Never one to shy away from confrontation, **Mike Saunders** discovers eight ways of communicating with outside world from the comfort of his own armchair...

Don't you just love confrontation? Asking someone out, telling the wife a divorce is imminent, or demanding a pay rise from your boss? Chances are you're like 98% of the population and can't bear it. Getting your feelings out face-to-face can sometimes be a monstrously terrifying experience.

Fortunately, though, the advent of such technologies as email and Internet conferencing has eased these ordeals somewhat. Instead of staring the bloke next door straight in the eye and interrogating him on his witnessed presence with your girlfriend, it's simple enough to put your finger to keyboard and set the issue straight.

Indeed, you may not gain the same satisfaction when gushing your feelings down the phone lines, but the unparalleled convenience and abundance of thinking time helps considerably. It's true that email (and newsgroups) get the most exposure of all network communication technologies, but the IRC chat system is still a very popular way to get 'talking' to people. Of course, while it's mightily useful for the situations mentioned above, IRC is also an

BACK IN GOOD OLD DAYS

Linux (the kernel) wasn't the only thing to originate in Finland. Way back in 1988, Jarkko Oikarinen created the first IRC client and server software. Initially used by friends and colleagues, and by the following year there were 40 servers in use worldwide.

Since then, the system has proved to be one of the most effective for holding online discussions and meeting up with friends around the world. A sizeable number of IRC networks have sprung up, including EFnet, Dalnet and IRCnet.

excellent system for general banter as well.

What's it all about?

So, why use this system when countless websites provide online chat rooms? Firstly, and perhaps most importantly, IRC doesn't rely on browsers, Java and hefty downloads, being an independent system for chatting over a network. Sites such as Yahoo! feature workable chat areas, but they provide all sorts of extra frills and fancies that most of us can easily live without.

When chatting with IRC, you simply fire-up a specialised program, point it at the relevant server, join a channel that takes your fancy, and get talking. There's no trawling through websites or making sure your browser is compatible – IRC is a stand-alone system. Anyone is free to develop a program that talks to IRC servers, and a variety of free chat clients have popped up for GNU/Linux and other UNIX systems.

As with email and newsgroups, there are certain conventions and unwritten rules that should be adhered to when you're using IRC. Naturally, chatters are expected to be polite, concise and able to tolerate others with opposing views. And in the event that someone becomes too abusive or starts



Pop over to <http://www.irc.org> for all sorts of IRC information.

adding unrelated nonsense – including blatant plugs –, the channel operators (or "ops") can remove that person from the discussion.

Fortunately, a "kicking" in IRC terms doesn't result in any broken ribs, but is simply a command performed by an op to throw the misbehaving user off the chat channel. Typically, people acquire op status when they've involved in running the channel, or are familiar and trustworthy to the head honchos in charge.

What features do I need?

As always, the requirements of users for this type of software vary greatly, and a newcomer to the world of IRC won't rely on some of the advanced goodies that old-hands use. However, because of its nature – pretty much constant typing – it's highly useful to

have a decent range of keybindings (and possibly user-configurable ones too) with tab-completion of nicknames and command-history recall (using the cursor keys).

Also, an intuitive and attractive interface is very much mandatory too. As you may be sitting at a three-hour discussion (with chatters joining and leaving all the time), it's essential that the display keeps you in touch with matters and doesn't deteriorate into a complicated mess.

In this issue's roundup, we're putting eight of the most prominent IRC chat clients for Linux under the spotlight. As always, we examine the crucial issues of performance, stability and ease-of-use, while also checking out the more advanced features that each one provides. So, put your feet up, grab a coffee, and let's see who comes out on top...

BitchX



BitchX 75p3 – Web: <http://www.bitchx.org>

As we've seen from our previous roundups of mail clients, newsreaders and even web browsers, there's still a lot of interest in the Linux community around text-based software. The X Window System may be running on the majority of desktop machines, but users on older hardware or those who just prefer the speed and simplicity of console programs are keeping up development in this area.

Without a doubt, the most widely-known text-based Linux IRC client is *BitchX*. Started by two guys known as "Trench" and "HappyCrappy" as a script for the popular ircII client, in 1994 it became a stand-alone program under the control of Colten Edwards and is still actively developed today. Due to its huge popularity, a similarly large number of scripts for the program have also been developed.

BitchX runs under a number of UNIX variants, along with Windows and MacOS. In terms of package sizes, you're looking at around 800K for a pre-built RPM, but it should also be on your distribution CDs somewhere. The only notable dependency is the ncurses terminal handling library, which is installed by default on most systems, and Tcl for some of the scripts.

Smack my BitchX up

Typically, *BitchX* is started from the prompt with an IRC server specified. What's most apparent from the opening screen is the splendid and slick visual design and colour scheme. Yes it is a console app, but the coders behind *BitchX* have made full use of extended character sets and a contrasting range of colours to create a client which compares well with its GUI-based counterparts.

In some respects, though, it's not immediately intuitive and the online

docs will have to be consulted before you can make any significant progress. This is mainly due to the restrictions of the text-mode program itself. As a result, the multi-window approach has had to be emulated in a slightly different fashion.

BitchX can either split the screen into two when viewing several channels, or make use of hidden windows (which are navigated using the ALT and number keys). Thanks to the well thought-out use of colour and graphics, though, *BitchX* maintains a consistent and easy-to-follow display and is a pleasure to use. Full tab-completion and command-recall are supported.

Meanwhile, general operation is very familiar for those used to IRC, and there's a text-entry bar for messages and commands along the bottom, together with some status info (channel name, current time etc.). New windows can be created with the `/window` command, and the online FAQ provides plenty of pointers and hints.



One of *BitchX*'s funky opening screens, which aren't even aware of the limitations of text-mode.

Going further

As a result of its extensive development history, *BitchX* has absorbed many advanced-user features over the years and is a favourite among experienced IRC fans. These include built in support for CDCC (and XDCC) chat and file transfers, a "toggle" menu for alternating the status of various features of the client (cloaking CTCP info, nick-completion etc.), and flood protection support.

Although it has no online documentation for scripting, the developers recommend reading through existing scripts to see what they're capable of. You'll find a large range of scripts (see <http://scripts.bitchx.com>). As with *X-Chat*, this makes the program a solid choice for those who find limitations with their current system and want to go further.

BitchX is supplied with good documentation for all the standard IRC commands, together with more in-depth user guides and concise information with the built-in functions. As it's based on the ircII client, the configuration details are stored both in `.ircrc` and `.bitchrc` files, where extra scripts and the like can also be added.

In conclusion, it's fair to say that *BitchX* is a dream come true for long-time IRC users. It's fast, stable and fantastic to look at.

LINUX FORMAT Verdict

Interface:	9/10
Features:	10/10
Documentation:	7/10
Performance:	8/10
Stability:	9/10

Supremely powerful and versatile, *BitchX* is the coolest way to chat. **Rating 9/10**

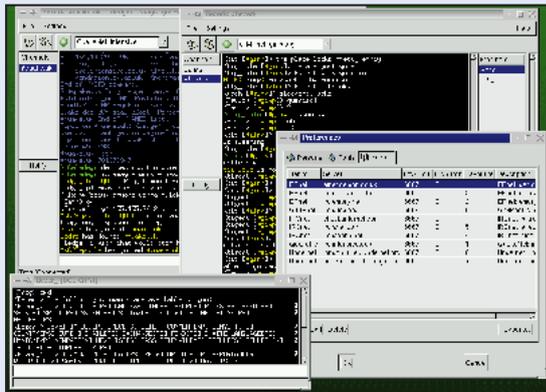
```

-#- BitchX: Join to #pcplus.test was synced in 1.408 secs!!
= Bud_Lite thinks he might try and write a blog in PHP
<M-Saunders> Zelenka: Lo
<Zelenka> M-Saunders: Howdy
-#- mode #pcplus.test [+o Zelenka] by tomgilder
<M-Saunders> What a star :)
<NBBhav> M-Saunders: she has a nice name. ;)
<tomgilder> IE is not "pretty crap"
<M-Saunders> NBBhav: Old Jan Dismas? :)
<tomgilder> it's a very good platform for HTML app development
<AlmostDevilish> ie rules all mankind
<Bud_Lite> It is quite good, IMO
<AlmostDevilish> Microsoft for President of computer land
-#- [Users(#pcplus.test:10)]
[ @AlmostDevi ] [ @Bud_Lite ] [ @LuRcH ] [ @M-Saunders ] [ @NBBhav ]
[ @Ng ] [ @techmoir ] [ @techroam ] [ @tomgilder ] [ @Zelenka ]
<AlmostDevilish> Bill Gates for the new God of the millenium
<techroam> tomgilder: if it can't match two colours which are the same...
<NBBhav> M-Saunders: I don't know, but it's a nice name. :)
-#- the_fuzz1121 [the_fuzz11@modem-103.selenium.dialup.pol.co.uk] has joined
#pcplus.test
User: Zelenka(+fiw) Server: adams.openprojects.net
Channel: @#pcplus.te(nt) Mail: 18 Time: 11:42pm Lag: 0
[ BitchX ]->

```

Once again, the conversation takes a downward turn as someone mentions Microsoft. Ho hum...

ROUNDUP | irc clients



Bezerk in action under the KDE desktop environment, displaying the very plain window furniture.

Bezerk

Bezerk 0.3.2 – Web: <http://www.gtk.org/~trog>

Software names within the Linux and UNIX communities are a weird and wonderful thing, and probably warrant a whole article on their origins and similarities. Just as mail clients tend to have a tree theme (first *Elm*, then *Pine*, *Spruce*, *Balsa* etc.), IRC software developers go one bizarre step further with *Zipper*, *BitchX* and now *Bezerk*.

Being a Gtk+ based client, *Bezerk* runs on a variety of UNIX flavours and is released under the GPL. Pre-built binaries are available at just 560k, and there are no major dependency requirements you're unlikely to have on your system. While it hasn't seen much development since mid-1998, it's acclaimed as being usable for everyday IRC chatting.

Bezerk starts with a simple configuration dialog, prompting you for your nickname. From there, you're thrown straight into the main window, which features very little in the way of icons or switches. Simply, there's a drop-down list of (user-definable) servers to connect to, together with vertical panes for the channels and member details.

The main text window makes basic use of colour to prevent the conversation and status messages amalgamating into a confusing mess, but there are no symbols or other

devices (as used by *X-Chat*, *BitchX* and *KVirc*) to keep things clear. Tab-completion of nicks is supported, along with command-recall for the text-entry line.

Bezerk has a few advanced goodies such as multiple server options and support for DCC chat, send and get, while the configuration dialog box (although very empty) includes the ability to alter the font and aliases. Tcl scripting is planned – although not yet fully implemented – but the supplied help files are devoid of any real info.

With this in mind, it's hard to recommend *Bezerk* as a trusty IRC client yet. For newcomers, it needs a more helpful interface, along with some decent documentation. For advanced users, it needs the scripting functions up and running. There's clearly potential in there, though, and we hope some keen hackers will transform it into a major competitor.

LINUX FORMAT Verdict

Interface:	3/10
Features:	4/10
Documentation:	3/10
Performance:	8/10
Stability:	6/10

Very basic and unhelpful to the newcomer, Bezerk still needs a lot of work. Rating 4/10

xlrc

xlrc 2.3.6 – Web: <http://www.linuxlots.com/~xirc/>

Right from the start of the project, KDE's developers intended their desktop environment to be fully internet-friendly. By providing a mail client, newsreader and browser, they aimed to provide a complete suite of tools, and numerous IRC clients have been written with this in mind.

Xlrc, like all KDE applications, is built around the Qt toolkit as created by Troll Tech. Its developers have aimed for ease-of-use while retaining a workable feature set. The download size for binary RPMs is around 500k, and it will run fine under other desktops and window managers too.

With its chunky buttons, lack of icons and default blue-on-white colour scheme, *xlrc* isn't very pleasant to look at initially. The user is presented with an editable list of servers (preconfigured with a huge range of the most popular ones – a nice touch), and from there another window opens to display the message of the day (MOTD) and other info.

Then there's another window for choosing the channel you wish to join, and yet another where the chat takes place. This approach is definitely love-or-hate – some will prefer saving screen real-estate by keeping each

matter in its own place, but others accustomed to different clients may find it frustrating to operate.

Xlrc features the usual DCC chat/send/get functions and notification lists, and it also uses a subset of the standard ircII commands along with a similar colour scheme to the Windows *mIRC* client. There's no nick completion or command history, but the decent help files are available.

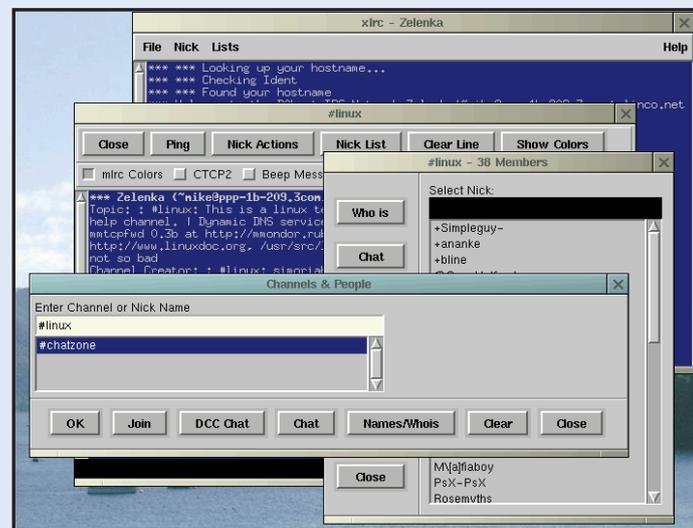
It's fair to say that *xlrc* is usable as a chat client, and the multiple-window approach may appeal to some, but others may find the setup limiting. Also, it doesn't support any form of scripting or plugins, and therefore it's not of great use to those with demanding requirements.

LINUX FORMAT Verdict

Interface:	4/10
Features:	3/10
Documentation:	7/10
Performance:	8/10
Stability:	7/10

A fast and simple client that's lacking in features, but the multiple-window system may be an attraction to some.

Rating 5/10



xlrc uses a multitude of windows, here showing the server, channel and people list boxes.

X-Chat

X-Chat 1.6.0 – Web: <http://www.xchat.org>

Perhaps the best known graphical IRC client for Linux, *X-Chat* has been in development for longer than many of its competitors in our roundup. Consequently, it has matured into one of the favourite chat programs for both newcomers and experienced users alike, and is released under the GNU GPL license.

One of *X-Chat's* main objectives is to run under a wide variety of platforms, and already it will compile under numerous UNIX variants (Linux, BSD, Solaris etc.) and even OS/2 and Windows. Most recent Linux distributions include a ready-built package, so it should be on your CD somewhere. If not, you'll be satisfied with the small download size of approximately 750k.

Dependencies include the *GNOME* and *Gtk+* libraries, which you're likely to have installed. Failing that, the project's site features a 'Lite' version which isn't so heavy on the extra package requirements. Although it uses bits of *GNOME*, it will happily run under other desktop environments like *KDE* and *Window Maker*, and can be built as a text-mode version too.

Ready, Steady, Chat!

ON launching, *X-Chat* offers a Server List window where you can choose

SCRIPTS 'N' PLUGINS

As one of its most notable features, *X-Chat* excels with its scripting and plugin support. Both the Perl and Python languages are supported, and many scripts are available on the main *X-Chat* website for things like timed kicks, operating external programs from the chat prompt, and even a voting script.

Undoubtedly, this feature is a boon for those demanding more advanced features in their client, but don't want to get their hands dirty in the actual program code. Various snippets of documentation are available on the web to guide coders through the script-creation progress, and above all it's an excellent addition to an already fine application.

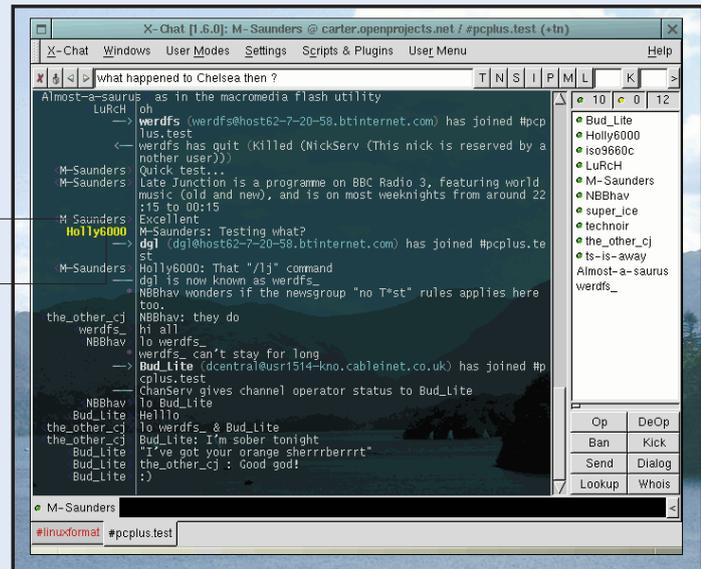
from a variety of IRC servers, featuring all the usual favourites. Also, this window permits the alteration of your nickname and personal details, which you can restrict if you want to preserve a certain amount of anonymity in your IRC activities.

This list is presented in a tree form, with

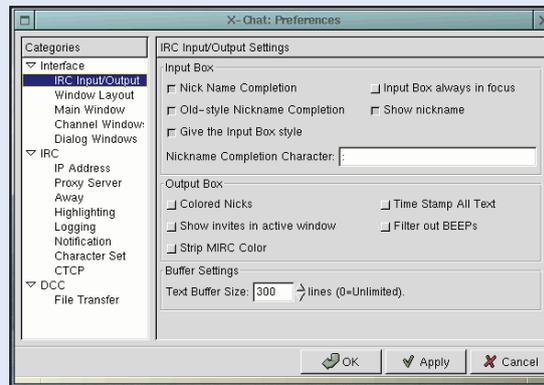
options for editing the various entries and skipping the MOTD if necessary. Connecting to a server places the user firmly in the channel window – which features the main pane for displaying text – and when a channel has been joined, a list of all the chatters resident there.

X-Chat sports a simple and obvious front-end, with the channel topic line along the top (beside some switches for modifying the status if you're an op) and the main text entry box sitting at the bottom, featuring tab-completion for nicknames. Each channel and server you join uses *Gtk+* tabs, which makes for quick working when you're talking in a number of places at once, and command-recall in the text-entry field is also included.

Despite its initially simple interface (which doesn't amaze in the



A busy evening in #pplus.test, with all sorts of heated debate and ramblings.



X-Chat's comprehensive Preferences box, with a huge number of options to tweak.

same hedonistic way as *BitchX*, for example), *X-Chat* is packed with features for tweaking its appearance to the maximum extent. Firstly, the main text window can be made transparent (with user-configurable tint) or have an image applied, while the font in use and word-wrap can also be tuned as well.

In all, it's clear that *X-Chat's* developers have worked hard on producing an interface which is straightforward and accessible to the new user, yet has enough potential for modification to keep experienced chatters content. It's a fine application all round.

Deeper discussion

Still, *X-Chat* isn't just a newbie's program to prepare them for things to come. The program astounds in the

huge number of advanced features and things that can be configured. For example, the main Preferences dialog box allows you to tweak virtually every aspect of the client, from basic colour-schemes and window layout options through to advanced proxy configuration options, character set translation and DCC file-transfer details.

As it stands, *X-Chat* takes the lead in the graphical IRC client race. Through a solid development history and clear focus on satisfying a wide range of users, *X-Chat's* programmers have produced the ideal introduction to the world of IRC for first timers, while providing all the extensibility and configuration necessities that experienced users demand. Definitely the best client if you're just getting started with IRC.

LINUX FORMAT Verdict

Interface:	10/10
Features:	10/10
Documentation:	7/10
Performance:	7/10
Stability:	8/10

A superb all-rounder, packed with features galore, and recommended for users of any experience.

Rating 9/10

Irssi

Irssi 0.7.96 – Web: <http://irssi.org>

While *BitchX* may indeed be the best-known text-mode IRC client, a decent number of other console alternatives do exist. *BitchX* has associated itself with its slick graphics and glamorous style, and has become enormously popular as a result, but of course many users want to avoid such superficial frills and are looking for something which just gets the job done without fuss.

Typically, many older programs have recently seen developers make progress on graphical front ends in an attempt to bring their software to a wider audience. But Irssi is different in that previous releases were built around the Gtk+ toolkit and the coders in charge have since opted to concentrate on text-mode.

Still, one of the goals of *Irssi* is to be modular and open to extension, so that other hackers can tack-on graphical user interfaces of their choice. Also, as a messaging system, *Irssi* is not entirely confined to IRC, and the projects leaders have expressed a wish to support other protocols like ICQ in the future.

We built the program from source, with the .tar.bz2 archive weighing in at around 550k – an easy download, and no doubt reduced since the GUI code was removed. RPM packages are also available, and the only significant requirement is the Perl language libraries, which are installed with most distros these days.

First impressions

What's most obvious from *Irssi*'s default visual configuration is the balance between aesthetics and functionality. Rather than go all-out on effects and colours (as with *BitchX*) the designers have been stricter with the display, providing a solid and comfortable appearance.

Along the top lies an inverted bar holding the topic line for the channel, then follows the main text pane, and at the base, a status bar containing the time, current nick, channel and email status. Beneath that is the text entry line, which supports tab-completion on names. Colour has been liberally used for status messages and highlighted names too.

Altogether, it resembles a less intense version of *BitchX*. Many of the keybindings – such as moving between hidden windows for channels and servers – are identical, and users with experience of other text-based clients won't have any difficulty getting started.

Looking closer

Irssi's programmers have put considerable effort into automating as many processes as possible, in an attempt to make the program rapid to work with. These include the automatic creation of windows when new servers and channels are joined, instant joining of specific channels after you've connected to a particular server, and more.

Similarly, much attention has been paid to the handling of multiple servers. You can connect to as many as necessary, and if a connection is dropped, Irssi can try again straight away. Also, servers are not tied to windows (as with ircll-based clients), so you don't lose the connection if you close that window.

It features the usual flood handling, logging and DCC chat/send/get. To handle its scripting capability, the program uses the popular Perl language which works

```
http://technoir.uk.net/
[23:48] @M-Saunders> Zelenka: Lo
[23:48] -! mode/#pcplus.test [+o AlmostDevilish ] by M-Saunders
[23:48] -! mode/#pcplus.test [+o the_fuzz1121 ] by M-Saunders
[23:48] @Bud_Lite> Zelenka : ?
[23:48] -! mode/#pcplus.test [+o Zelenka ] by M-Saunders
[23:48] @AlmostDevilish> (sounds like a Russian name)
[23:48] @Bud_Lite> Zelenka : ditto what?
[23:48] @Zelenka> AlmostDevilish: Nah, I'm Czech.
[23:48] @AlmostDevilish> Oooh
[23:48] @AlmostDevilish> Fellow European =o) I come from nearby to y ou
[23:49] @Zelenka> AlmostDevilish: Yeah, but I died 256 years ago.
[23:49] @M-Saunders> Zelenka: Lo
[23:49] -! the_fuzz1121 [the_fuzz11@modem-103.selenium.dialup.pol.co.uk] has
quit IRC [xchat]
[23:50] @M-Saunders>
[23:50] @M-Saunders>
[23:50] @M-Saunders>
[23:50] @M-Saunders>
[23:50] @M-Saunders>
[23:50] @M-Saunders>
[23:50] @Bud_Lite> oh no.
[23:50] [Zelenka(+fi)] [Z.#pcplus.test(=nt)] [Act: 1] [Mail: 17]
[!#pcplus.test]
```

Irssi's main chat window, with the topic along the top and status line near the bottom.

using a system of signals that are comprehensively documented in a dedicated help file. The remaining documentation is thorough and well-written, and improves the application's standing as a mature, stable and advanced client.

Irssi also claims to being one of the most secure IRC clients available today, with the lead hacker pointing out that none of the "ugly kludgy" code from ircll (which many other clients are based on) exists in the source code.

The final judgement

In the end, *Irssi* makes an excellent alternative to *BitchX* for those running at the text console. While it's not as fancy to look at, the careful design and fine documentation put it up there with the best. Being one of the

fastest clients in our roundup, Irssi also proved to be very stable and solid to use.

Of course, it's still not the easiest route into IRC for newcomers, and the GUI-based programs will undoubtedly be a better introduction for those unfamiliar with the chat system. However, if you're looking for a clean and fast console tool without any fancy trimmings, take a look.

LINUX FORMAT Verdict

Interface:	8/10
Features:	8/10
Documentation:	9/10
Performance:	10/10
Stability:	9/10

Swift and robust, Irssi makes a highly competent alternative to BitchX. **Rating 9/10**

```
#pcplus.test "also look at http://www.re-con.co.uk and tell me why it won't wo
[00:11] -! Geminiani [nike@du-016-0060.claranet.co.uk] has joined #pcplus.test
[00:11] -! Topic for #pcplus.test: #pcplus.test "also look at
http://www.re-con.co.uk and tell me why it won't work"
[00:11] -! Topic set by AlmostHuman [Mon Nov 27 23:59:51 2000]
[00:11] -! Users(#pcplus.test)
[00:11] @Bud_Lite | @Ng | Geminiani |
[00:11] @M-Saunders | @technoir | | tomgilder |
[00:11] @NBBhav | @VortX |
[00:11] -! #pcplus.test: Total of 8 nicks (6 ops, 0 voices, 2 normal)
[00:11] -! Channel #pcplus.test created Sat Oct 7 00:23:17 2000
[00:11] -! Irssi: Join to #pcplus.test was synced in 2 secs
[00:11] < Geminiani> Hello again
[00:11] @VortX> lo Geminiani
[00:12] @NBBhav> lo Geminiani
[00:12] < Geminiani> Good to be here :)
[00:12] @NBBhav> Geminiani: you listening to R3?
[00:12] < Geminiani> Yep indeed.
[00:12] @VortX> er.... good to have you here :-)
```

Heavy discussion continues as the regulars desperately think of something profound to declare.

Zipper

Zipper 1.631 – <http://zipper.sybermedia.homeip.net>

Although **Gtk+** and **Qt** have become the two most popular user-interface toolkits for building graphical applications in recent years, those with long-term experience in the Linux and UNIX world know that there are others available.

One such system is the combination of **Tcl/Tk**, which provides a scripting language and toolkit for creating X-based software, and, as it is interpreted, it has become a favourite for those writing small programs. The language pops up in a wide number of tools, including the file manager *Tkdesk*.

Zipper is an IRC client built around **Tcl/Tk**, and the main developer (Ashley Bowers) started the project to get himself familiar with the scripting language. It was not designed to compete with the existing major players in this field, but simply to provide an alternative that satisfied developers' needs.

If you get hold of the `.tar.gz` archive, you're required to run a specialised installation script (which

also sets up various configuration options and similar details). Unfortunately, this system isn't very clean and involves modifying the path, but easier RPM packages should be available on the Internet. The most important requirement is version eight of the **Tcl/Tk** libraries, which are freely available.

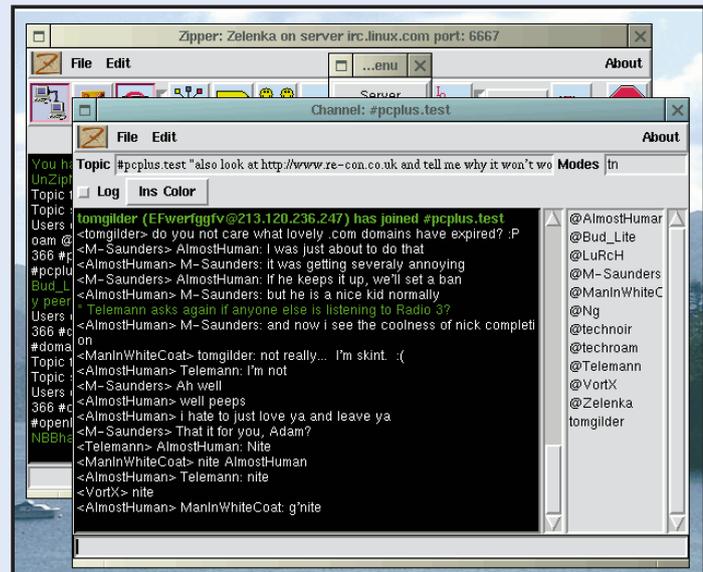
Up and running

Zipper's interface resembles many of its other graphical cousins, with a row of chunky icons along the top providing access to common operations, together with a status bar

In the end, *Zipper* doesn't really cut the mustard when compared to *X-Chat* and *BitchX*

(providing help for the buttons) and the essential text-entry field sitting along the bottom. Also, a separate window holds links to every channel window and private message you're engaged in.

The main text pane, where conversation takes place, is equally



Zipper in action, showing the main server window and a channel box.

sparse and features little in the way of colour or fancy effects. There's no highlighting when you're directly addressed, nor any indentation of nicknames. In fact, there's very little of

note at all, which doesn't help to make things clear when you're involved in a busy debate.

Zipper's default sans-serif font

renders the obligatory Figlet quips useless, and tab-completion hasn't been implemented either. While the latter is hardly a vital issue, it can be a godsend in lively debates when you're trying to catch someone's attention.

To be fair, the interface does an adequate job for general light banter, but if you frequent busy channels – and have a lot to say – it's simply not rapid or elegant enough to warrant constant use. Some tidying of the text layout and inclusion of essential keyboard shortcuts would improve this situation somewhat, but for now it's not a thrill to work with.

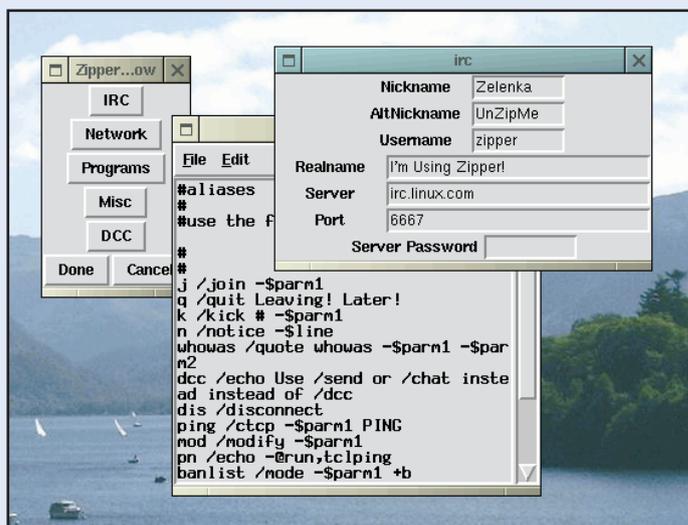
Other issues

Perhaps *Zipper's* most sorely-missed feature though, is the ability to connect to multiple servers. Again, this won't be a problem for casual chatting and for those who just pop into a channel now and again, but for regular IRCers it's something that will be very much missed.

Thankfully, *Zipper* improves its standing in places with a few other notable features. It sports **DCC** support and has a comprehensive range of configuration windows covering the key aspects of IRC setup. The fact that some of these just throw up a text box for editing config files is a slight let-down, but there's some fairly decent online help which explains the various features of the program and there's even some elementary scripting support with **Tcl**.

In the end, *Zipper* doesn't really cut the mustard when compared directly to the giants of *X-Chat* and *BitchX*. However, we have to bear in mind that the developer wasn't aiming for this, and many light users may find it a straightforward and speedy client to use.

Sadly, the code is no longer being maintained by its original author and it needs some extra inclusions before it will convert a lot of people.



Configuration of Zipper is split into several dialog boxes.

LINUX FORMAT Verdict

Interface:	4/10
Features:	4/10
Documentation:	6/10
Performance:	9/10
Stability:	5/10

An average all-rounder with little to turn the heads of experienced IRC users.

Rating 5/10

Conclusion

Anyone who has seen the growth of Linux as a significant desktop OS over the past few years will no doubt have noticed the increasing trend towards style and aesthetics. Software such as *Enlightenment*, *XMMS* and *Mozilla* (with its Modern skin) show a growing tendency to focus on appearance and design – but, of course, not at the expense of functionality.

However, the explanation behind this is very clear: these programs are used day-in day-out, and when you're staring at something for several hours on end, it has to be non-distracting, allowing you to keep your mind on the subject in question.

In this light, it's easy to see why the developers behind *X-Chat* and *BitchX* have spent considerable time tweaking their user interfaces. *X-Chat* features basic themeing, transparent tinted backgrounds and a neat layout. *BitchX* sports clever use of ASCII text

mode and colours, and peppers the chat windows with pretty icons and similar frills.

Together, these help to create a working environment that's fast to operate and easy to look at, which is vital when multiple-hours of eye-strain-free discussion are taking place. Huge, chunky icons, a garish colour scheme and fiddly controls would be a recipe for disaster in this

X-Chat takes the honours in our roundup – it manages to be an excellent client for all kinds of users

area, and the coders of these two main programs have demonstrated a very good understanding of the issue.

Which should I use?

As the results show, *X-Chat* and *BitchX* are currently the two main contenders for the king of IRC crown.

They've both had lengthy development histories, have both incorporated advanced features and scripting to satisfy experience users, and they've both had solid work on their front-ends to create clients that eventually become transparent to use.

If you've yet to experience IRC and want an accessible introduction, *X-Chat* and *KVirc* are your best choices (and are both included on most distribution CDs). With their helpful setup screens, abundance of tooltips and straightforward interfaces, either would make the ideal client to get started in the world of IRC.

It should be emphasised that *X-Chat* also performs well when it comes to more demanding requirements, and the capable scripting and plugin ability warrants closer inspection for experienced users too. It's for this reason that *X-Chat* takes the honours in our

roundup – it manages to be an excellent client for all kinds of users.

However, if you're not graphically minded, *BitchX* is right up there too in the features and usability stakes, and if you're not running X or fancy a change from the world of point 'n' click, it's well worth checking out. Because of their near-constant typing nature, it's worth noting that the text-based and graphical clients don't differ greatly in use, and you may be pleasantly surprised at the speed and ease of working at the console. 

CHAT HERE TO BEGIN

One LXF fan has created an Unofficial Linux Format IRC channel, where readers of the magazine can go to ask questions, talk about the mag, and have a natter with the other regulars. To take part, simply use the following details in your favourite IRC client:

Server: `irc.linux.com` **Channel:** `#linuxformat`

Also, check out the `#pcplus.test` channel on the same server for even more off-topic banter.

Look out for such regulars as NBBhav, `the_other_cj`, `Bud_Lite`, `iso9660c`, `Stan103`, `LuRcH`, and of course myself (M-Saunders). Speak to these people at your own risk, though, as there's no guarantee you'll get a remotely sane response...

TABLE OF FEATURES

We tested the IRC clients on a 800MHz, 320Mb RAM PC running Linux-Mandrake 7.1. Note: the package sizes and memory usage statistics below should only be used as a rough guide. Memory usage varies from system to system, and is affected by the program's configuration and various other factors.

	BitchX	Bezerk	xirc	X-Chat	Zipper	Irssi	EPIC4	KVirc
License:	GPL	GPL	GPL	GPL	GPL	GPL	Distributable	GPL
Toolkit:	Text	Gtk+	Qt	Gtk+	Tk	Text	Text	Qt
Dependencies:	ncurses	Gtk+	Qt	GNOME, Perl	Tcl/Tk	Perl	ncurses	kdelibs, Qt
Package size:	850k	560k	505k	730k	230k	900k	650k	1850k
Memory usage:	2500k	1800k	6500k	5300k	3700k	1604k	1500k	8500k
Completion/recall:	Yes	Yes	No	Yes	No	Yes	Yes	Yes
Multiple servers:	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
Scripting:	Yes	No	No	Yes (Perl, etc.)	Basic support	Yes (Perl etc.)	Yes	Yes

THE BEST NEW OPEN SOURCE SOFTWARE ON THE PLANET

HOT PICKS



All that's new, exciting and open source is given a good grilling by Jon Kent.

Service is almost back to normal on the CD this issue. Due to a few last minute hitches, we don't have all the HotPicks candidates for you, but we're getting there.

There's a little bit of a web feel to this month's hottest stuff, with

HTML editors, proxy servers and a network analyser. More proof, if any be needed, of the increasing dominance of Linux in the internet sphere. If you have ideas for things that should be Hot Picks, drop us a line!

Nautilus

Version: 0.5 (Preview release) **Web:** www.eazel.com

Nautilus is a file manager (or Desktop Shell as Eazel – the developers – like to refer to it) on steroids, that been designed to work on the HelixCode *GNOME* desktop. Currently it is available as a preview release (version 0.5 at the time of writing), it is therefore a little bit rough around the edges, but is showing a lot of promise.

Installing *Nautilus* is fairly straightforward on a RPM based system, as Eazel provide an installation utility which can be downloaded from the web site. For DEB based systems, there are packages available, so update your `/etc/apt/sources.list` file with the following entry to install:

```
deb
http://people.debian.org/
~kitame/gnome/release ./
```

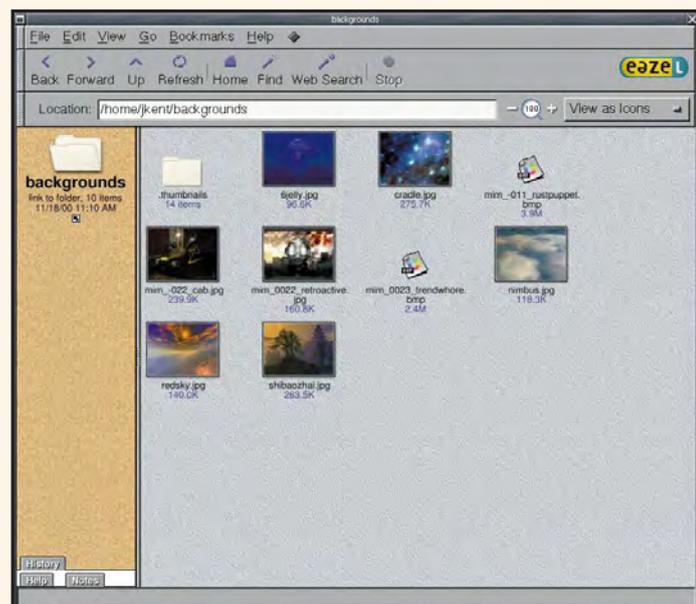
Then run an `apt-get update` and `apt-get install nautilus` to get the packages installed on your system.

When you run *Nautilus* for the first time, you are presented with a setup window that asks to indicate the type of user you wish to be

considered as, the options being, beginner, advanced or hacker. What you choose to be regarded as has a minor impact on how *Nautilus* presents data to you, such as whether or not you can see hidden files or not. It's an intriguing, if vaguely pointless, feature and it will be interesting to see if reaction to it will be reflected in more in-depth changes in the full release.

Once you have identified the type of user you wish to be, you are then presented with a window that reminds you that this is development software. Once you've closed this window you're ready to go. This warning window appears every time you start up *Nautilus* which can be a bit annoying after a while. It would be nice to be able to switch it off but there is no option for that.

With *Nautilus* up and running you are presented with a fairly standard file manager display, with the contents of your directory in the right hand panel, and the directory overview on the left panel. However, with the menu icons there are options not usually available within a



Graphic files viewed as thumbnails. Do we really need to tell you how useful this feature is?

standard file manager. The Web Search icon points to the fact that this is not just a file manager, but can also be used to surf the web. Therefore entering a web site within the location bar will display that web site, with all the usual support. In fact *Nautilus* uses *Mozilla* via the *Mozilla* application framework to provide this function, therefore you will need *Mozilla* installed to make use of it.

The other option that is unusual is the View As drop down option list. This provides some very neat, useful features. If you go to a directory

which contains pictures, *Nautilus* will display small but distinct versions of these images, allowing you to quickly find the file you were looking for. Double clicking on the image causes the program to display a full size version of it. Using the View As feature you can change how graphic files are displayed, so if you just want a list of files you would select View As List. Another example of the powerful nature of *Nautilus* is the ability to play mp3 files. When you go to a directory containing mp3 files the view changes to View As Music, and you

can simply select an mp3 to be played. The available views are dictated by the contents of the directory you are browsing, so if you have images and HTML files, these views are added to the options.

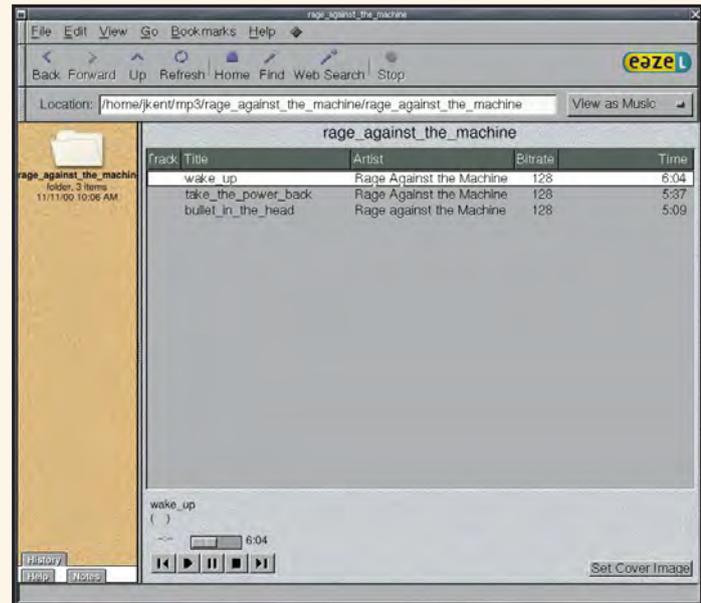
Nautilus itself is highly configurable, although if you have set your user status to 'beginner' you are limited to modifying just the look and feel of the desktop and not how file and directories are displayed. *Nautilus* supports themes, and has various backgrounds and colours that can be applied to each panel by simply dragging and dropping the image or colour onto the panel you want to alter. You can even add labels (or emblems) to files giving extra descriptions of the file, such as Important, New and OK.

When you add directory specific changes, such as background changes, *Nautilus* writes a file called *nautilus-metafile.xml*, which is an

XML file describing the changes you have made. Therefore when you re-enter this directory this file is read, if it exists, and the changes are implemented. This is an interesting idea, as most applications that perform this type of function will write any display details back to a single file.

As this is still a preview you should expect some program crashes to occur. *Nautilus* development is looking very good for the future; it is expected to replace the existing default file manager in the next release of *GNOME* and is more user friendly than its predecessor ever was. It's fairly stable and, occasional crashes aside, is also very usable.

Along with *HelixCode*, *Nautilus* goes some way in helping to make the *GNOME* environment more polished and user-friendly. Overall, this is a very eye-catching application, with some interesting features.



A good example of the context sensitivity of *Nautilus*: a directory of mp3s can be played straight from the desktop.

Whether you feel that all these bells and whistles make sense as part of a desktop file manager or not, you

will be impressed by the way *Nautilus* functions and its suave good looks. Get it now.

Amaya

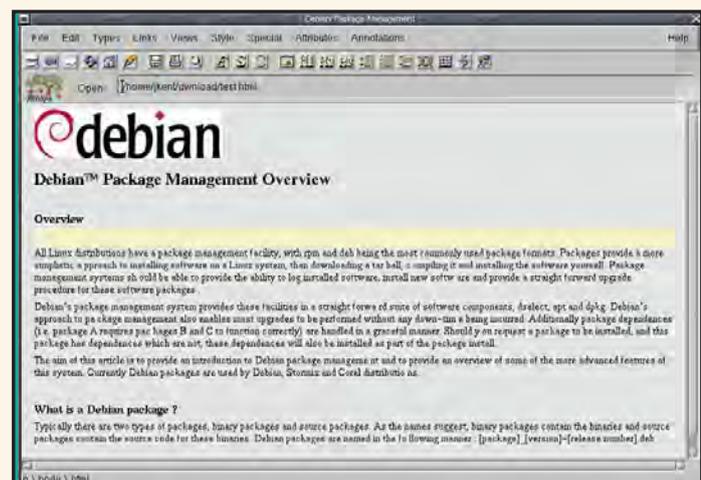
Version: 4.0 **Web:** www.w3.org/Amaya

Amaya is a browser and a WYSIWYG web page authoring tool. The current version implements HTML, XHTML, MathML, SVG, CSS (Cascading Style Sheets), and HTTP. As it is created by the web standards group, *Amaya* complies with all the current HTML standard, but does not support the extensions that Netscape and a certain company from Redmond, Washington have bolted on. It's actually very interesting to surf the web using the browser function of *Amaya* to see which sites do not conform to the letter of these standards – RedHat.com for example fails to load correctly. Whoops! This rigid conformity to W3C standards means that any web pages that are produced using this product will work across all browsers, including console based browsers such as Lynx. To surf the web using *Amaya*, you need to type in the full HTTP location details

in the Open dialog, such as <http://slashdot.org/>, otherwise it assumes that you want to open a file held locally on your system.

As is expected of a WYSIWYG web editor, it is very easy to jump straight into creating a web page with *Amaya*. The icon tool bar provides most of the common tags and functions that would be required with additional tags available from the standard menu. To start creating a web page you simply select File and New XHTML, MathML, CVG or CSS and a new window is open. You can now start typing the text of the page, adding the usual text mark ups as they are required.

There is support for cascading style sheets (CSS) and an external CSS file can be linked to the document you are creating. You can also place these style markups straight into your HTML page. These



WYSIWYG page authoring lets you see exactly what your viewer will see.

functions are tucked-away under the Style menu, allowing you modify the usual markups such as text font and size. Other HTML markups such as tables, forms and lists are provided, either from the icon tool bar or from the text tool bar.

Amaya supports the annotation of comments or explanations that can be attached to the page being created. This allows you to effectively comment on your page for future reference. Annotations do not form part of the page itself, but are

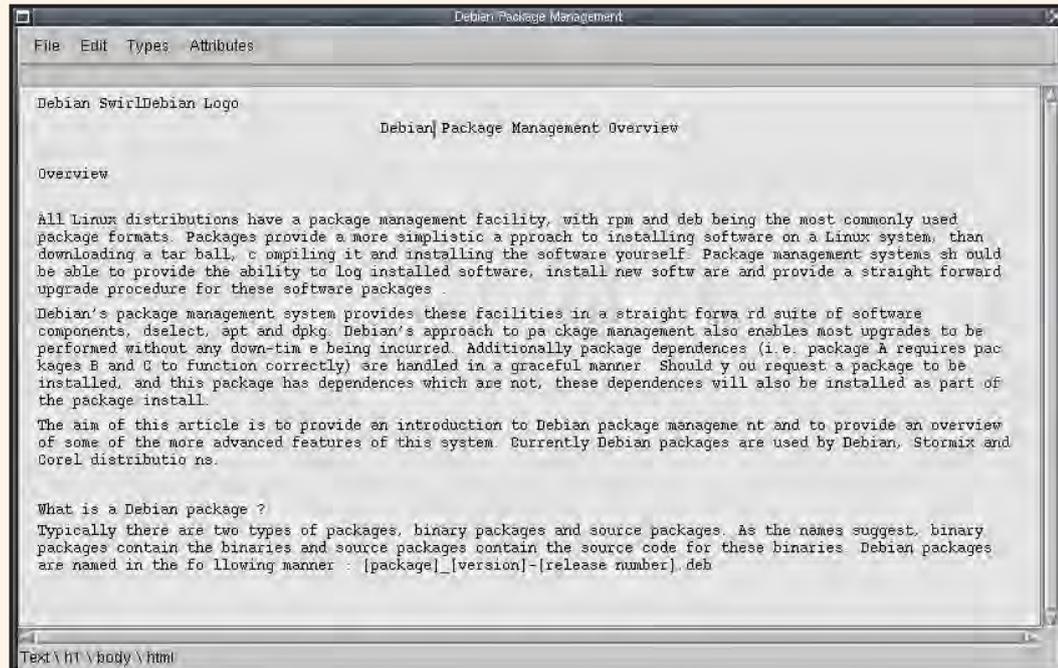
external documents that *Amaya* references when you edit the page (assuming that an Annotation file reference to your page exists). You can either add an annotation to the whole page or just a section. This is a forward-looking feature of the software that could prove to be very useful. You could use the usual <!-- [comment] --> structure to add comments to your page, however, this would mean editing the HTML directly which defeats the object of using a WYSIWYG editor in the first place. →

REVIEWS | hot picks

→ *Amaya* will insist that you supply certain information to ensure that your page conforms to the W3C standards. A good example is when you insert a graphic into the page. The program will refuse to insert the graphic unless you supply text for the 'alt' entry that should be present next to a graphic file. This entry allows console based browsers to display a meaningful text message instead of the graphic itself, therefore making the page more accessible across all types of browser.

Of course, no self respecting web authoring tool is complete without being able to modify the underlying code directly, and *Amaya* is no exception. In fact, it has a few unusual views in addition to viewing the HTML code. For example the view Show Alternative displays the page as if you were accessing it from a console browser, and the Show Table of Contents displays a window with all the heading listings in order.

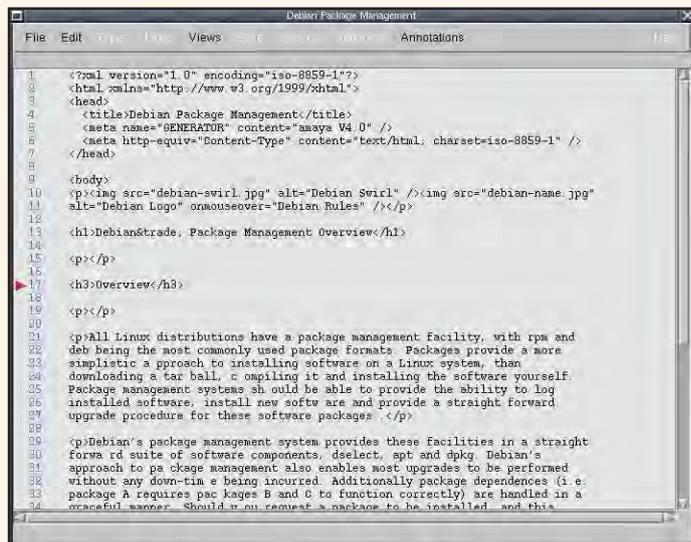
As this software has been designed to allow you to create a web page using the latest standards, you need to be aware that some of these standards have not been fully adopted in all browsers available. Indeed, some of the technology in *Amaya* is at the leading edge of web development, therefore older browsers, such as the Netscape 4.7.x series, will have problems with functions such as the Maths symbols that can be added. However, Mozilla



Amaya's view pages as console feature lets you tailor your code for every type of user.

(build M18) appears to work correctly with pages generated by the package.

This is a stable and very impressive WYSIWYG web authoring tool, certainly one of the best available for Linux. With its firm adherence to W3C standards you can be confident that everyone can view your page with any browser. There is no need for the "Best viewed with Netscape/Explorer" rubbish that is common on the web, as *Amaya* ensures that you do not fall into this application specific trap. In short, if WYSIWYG is your preferred method of authoring you can't go far wrong with *Amaya*.



Don't like WYSIWYG? You can also go in and edit your HTML code by hand.

wwwoffle

Version: 2.6

Web: <http://www.gedanken.demon.co.uk/wwwoffle/>

Wwwoffle is a proxy server with several interesting additions. The most interesting feature of *wwwoffle* is the ability to cache pre-defined web sites so that they can be viewed whilst not being connected to the web. It also has the intelligence to only update this local cache should the page(s) change or a predefined timeout be reached. Much like the *wget* utility, *wwwoffle* recursively fetches these pages down to a specified depth and has the ability to follow links contained within these web pages. This local caching can therefore be used to grab your favourite web sites whilst you surf elsewhere or check your e-mail. The cached web sites can then be viewed off-line, therefore reducing your monthly phone bill.

The proxy server function of *wwwoffle* provides the usual facilities, such as site blocking, and page replacement of a blocked site

(commonly used to display a warning message). You can also configure *wwwoffle* to auto-connect and collect pages that are not currently held on your hard disk.

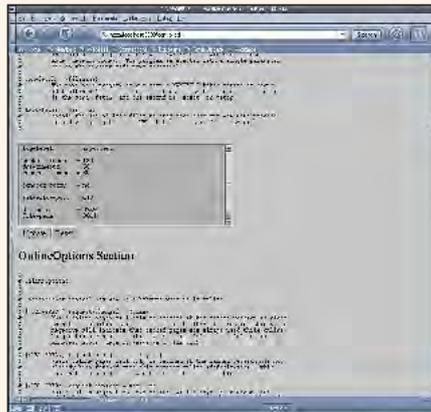
The main configuration of the software is controlled by the configuration file *wwwoffle.conf*. However, configuration of web sites to be cached, and other high level controls, are set via a web interface. The web interface also provides the ability to edit the configuration file. *Wwwoffle.conf* is a fairly large file, weighing in at just over 1000 lines – including comments – with a default configuration. This is well commented throughout making it easy to understand the application can be modified to meet your needs. You should not have to make many changes to the configuration, but there are a few things that must be altered out to ensure some of the

defaults are correct. These include the renumbering of the default HTTP ports used and adding a username and password to control access to the web based administration tool.

To access the web page configuration screen, you point your browser to the server running *wwwoffle* and the port it is bound to, for example:

`http://localhost:8080`. From this main page you can access all the major functions. For example, to add a new site to be cached locally, you would select the monitor-options function under the section WWWOFFLE Monitor Page Requests. On this page you enter the fully qualified name of the web site you wish to cache/monitor (for example `www.uk.debian.org`) and apply month, day and time conditions to control when this page is updated if required. Additional conditions can then be applied to the web site, such as the depth required, how links should be treated, should images be fetched etc. This is configured via the `/refresh-options/` option.

To use *wwwoffle*, you need to configure your browser to use it as a proxy server. With *Netscape* and

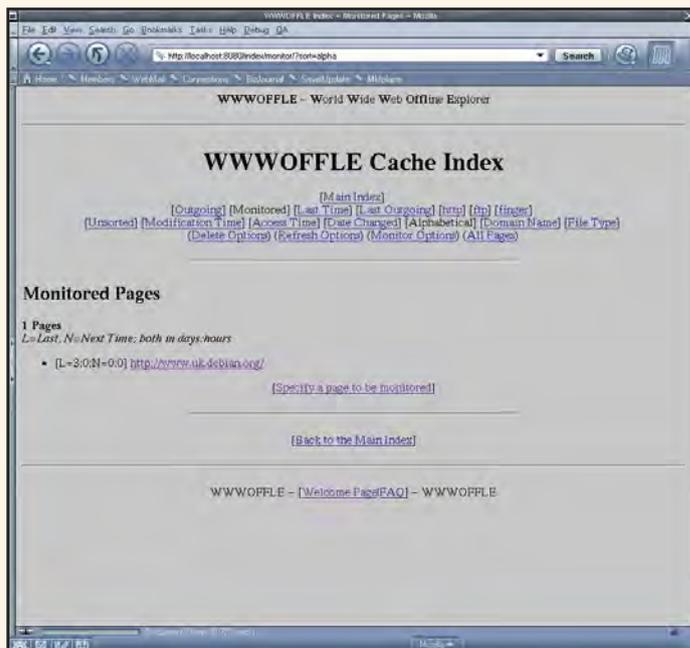


You can set up the software using the excellent web-based front end.

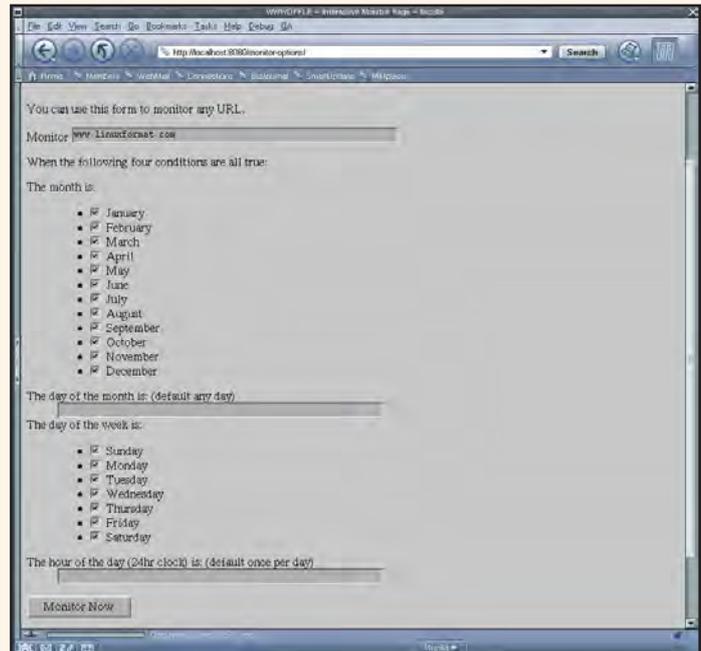
Mozilla this is configured via the Edit – Preferences menu option, under Advanced. By default *wwwoffle* uses port 8081, so you would configure this port against the protocols you wish to use (i.e. HTTP).

Wwwoffle works well as a basic proxy, but the additional functions make it a valuable, especially its ability to cache web sites locally for off-line viewing. Configuration is straight forward, as is the maintenance of the program. For home use, the caching function could be a god-send, because of the potential savings you could make.

Overall this is a well thought out and implemented product which is stable and has no major omissions.



Wwwoffle can grab your favourite sites for off-line reading while you do something else, like check your email or make some coffee.



Setting up new caches is easy, then just sit back as the software does your bidding, saving you time and money.

ethereal

Version: 0.8.14 Web: www.ethereal.com

Ethereal is a GUI network analyser which allows you to interactively browse network packet information from either a live network or from a capture file. The program understands network capture files that were generated by other network analysers such as *Snoop* and *tcpdump*. *Ethereal* needs to be run as root in order to capture all packets on the network that you wish to analysis.

When the program is first launched, you are presented with three panels displaying a summary of the packets received, a detailed description of each one and a hex display of their contents. Highlighting a packet within the summary panel will then display the detailed information and hex display of it in the associated panel. Within the detailed description you can drill down into the packet by selecting the '+' next to the option of interest. If you highlight one of the lines within the detailed summary, this area of the

packet is then displayed within the hex display panel. *Ethereal* can also assemble packets within a conversation together in a window by selecting Tools – Follow TCP Stream.

Usefully, you can apply filters to ensure that you only view the network traffic you're interested in. These filters can range from the very simple to the quite complex. *Ethereal* filters allow the use of standard symbols, such as `==` or `>=`, along with logical expressions, such as `&&` or `!`. The use of these common symbols and expressions allows you to create the logic for the filters quickly. However, before you can write any filters you need to understand how to reference the internal filters that are supplied with *Ethereal*.

Ethereal supports a wide range of protocols, and fields contained within these protocols, which you need to use for filtering against. This can be both useful and a mild hindrance, as you will probably →

REVIEWS | hot picks

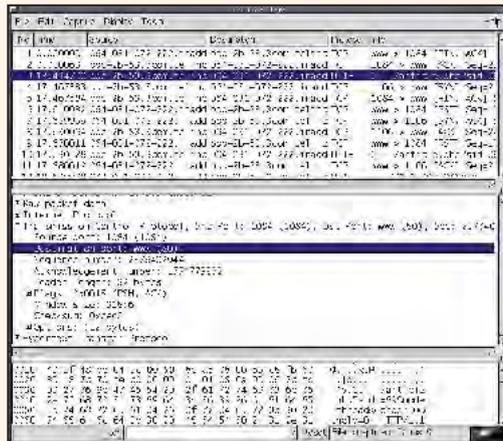
➔ have to refer to the man page each time you want to filter a different protocol. It does provide cleaner analysis, though. A very simple example of a *Ethereal* filter is:

```
ip.addr eq 172.1.243.59 &&
tcp.port == 23
```

This filter would only display telnet packets if the ip address is matched (that could be configured using a capture filter). With *Snoop* the rule would be written as:

```
snoop host 172.1.243.59 port 23
```

This is perhaps a little cleaner for a simple analysis, however *Ethereal's* filters display their power with complex filters. By default the app will capture all packets on the network



Ethereal can perform a comprehensive analysis of your network.

Aside from standard network analysis, *Ethereal* also lends itself to several other uses. Because of the configurability of filters, *Ethereal* could be used as

being listened to. The data being capture can be modified at start up (using **-f [capture filter]**) or when the capture is started by adding the filter within the Filter Input window. The data capture filter that can be defined at startup following the same syntax as *tcpdump*.

an Intruder Detection System (IDS), by configuring the filters to look to traffic that it considers suspicious.

Ethereal has a reasonable complete set of preferences for the basic functions of the GUI. You can modify the columns of the various TCP streams, modify locations of

scrollbars, and change some basic printer settings. Additionally, the capture display options can be modified slightly.

Unfortunately there is not much documentation, aside from the man page, and there are few examples that you can learn from. This seems a shame as *Ethereal* is a very impressive product. Stability wise it's a little temperamental, crashing from time to time – especially when looking at captured data. But this is to be expected from a product that is still in the development phase.

Overall *Ethereal* is an impressive, usable product. Because of its pre-release status, its use for critical work is not to be recommended at present. Once the stability issues have been resolved though, *Ethereal* should prove to be an indispensable tool.

GNOME-APT

Web: www.debian.org

As outlined previously in *Linux Format*, Debian's apt-get command line utility is an extremely powerful package management utility. *GNOME-APT* is a GUI front end that provides the same functionality as apt-get, allowing you to add and remove packages easily. It uses exactly the same methods to retrieve packages, i.e. http, ftp and file access types.

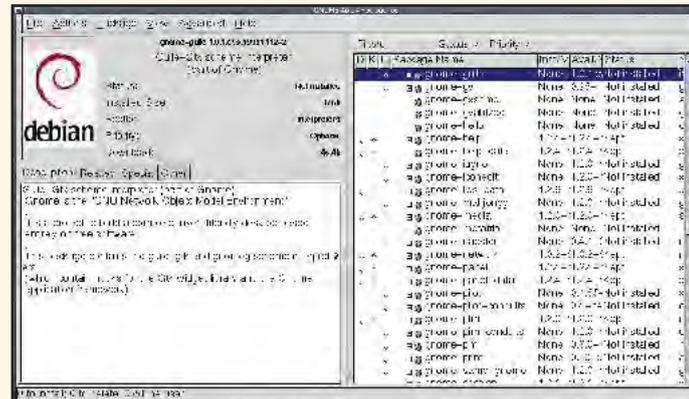
When the program is started you are presented with a single window containing three panes. On the right is the package control pane, the upper pane on the right displays a summary of any package that is highlighted, and the lower right one provides a more in-depth description of the package selected. Selecting a package to be installed is actioned by selecting the radio button under the I (Install) column, while to remove a package you select the radio button in the D column. Similarly a package can be put on hold by selecting the radio button in the K (Keep) column. The name Keep is strange as



GNOME-APT gives you a list of searchable packages

packages in this state would normally be referred to as 'on-hold' in Debian package speak.

There are some useful features that make *GNOME-APT* more than just a run-of-the-mill apt-get front end. There is a handy search facility available that will pop-up an additional window listing all packages which meet your search criteria. If you select one of the results, this



The search facility gives GNOME-APT an edge in the management stakes.

package is then displayed in the various panes of the main window. Additionally you can directly edit the sources.list file from within the program and run an update from the source listed. Installation and upgrade can be performed separately or together, which gives the user much more control.

The GUI itself is quite configurable, allowing you to modify how the columns appear and which are shown. How the packages are grouped together can be changed from standard dselect view of showing packages by section. In addition to displaying the package summary, icons are used to identify conflicts, dependencies, recommendations, suggestions and

replacement of the package. This is a fairly useful feature, however no more detail is supplied to aid you, such as the packages that are depended upon. This is a shame, as this would have made *GNOME-APT* an even better program.

Ultimately whether you will find this program useful will depend on how up to speed you are on apt-get and whether you prefer a GUI to the command line. It does offer some advantages, most notably the search function, which might prove very useful. It also benefits from being a stable application that can be used without fear of falling over half way through your job. So if GUI is king, and you have a Debian based system, this could be your best friend.

FEATURE | 2.4 kernel

THE 2.4 KERNEL FOR NEWBIES

If you are new to Linux, then you will probably fall into one of two categories: a home user running (or thinking of doing so) Linux on a desktop or laptop PC, or a business user running Linux on a server or desktop machines, that are not necessarily x86 based.

Many of the issues discussed in this article will be a little irrelevant to home users: for example the Linux port to the IBM S/390 mainframe. However, despite this, there are many innovations in the 2.4

kernel that will be of importance to mainstream desktop users. Some examples of this are the new and improved support for USB and Firewire peripherals.

In some ways, there are more goodies in store in 2.4 for business and corporate users than for home users. Some examples of this are kHTTpd, support for files greater than 2GB in size, and an increase in the maximum number of users per machine from about 65,000 (2^{16}), to 4.3 billion (2^{32}) in 2.4.

UNDER THE SKIN OF 2.4

The 2.4 kernel promises to make your Linux system fitter, happier and more productive. But what's it really like?

Chris Howells is our man with the scalpel...

New Features in the 2.4 kernel

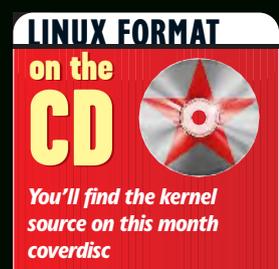
The name 'Linux' refers to the software which is right at the heart of your Linux operating system. Some distributions insist that the proper name for the entire system is GNU/Linux (acknowledging the GNU heritage) because the Linux kernel is such a small part of the entire system. However, the kernel is an extremely important piece of software, and amongst other things, provides the ability for your applications to interact with the hardware. It would be no good to have all the software such as the GNU tools, if you didn't have an operating system kernel to run them in!

The release of the 2.4 kernels is based on work added to the 2.2 kernels. The

series of development kernel between the 2.2 releases and the 2.4 release were, not surprisingly, labelled 2.3. Currently, the unstable development status has been upgraded to 2.4.0test – better than the 2.3 development kernels, but not good enough to enter a 2.4.0 stable release. This article is based on the latest 2.4.0test kernels, although throughout, we have abbreviated this to simply 2.4, in anticipation of the similar feature set in 2.4.0test and 2.4.0 final.

Why is the kernel so important?

As the kernel is such an important part of the operating system, the release of the Linux kernel is one of the more important events in the calendar of the Linux community. Since development began on the Linux kernel in the early 1990s, the number of features provided by it has steadily increased, as has the size! For comparison purposes, release 0.01 of the



FEATURE | 2.4 kernel

Web resources

<http://www.ia64linux.org>
– Information about running Linux 2.4 on Ia64 (Itanium).

<http://www.ibm.co.uk/linux> – Information about running Linux on IBM hardware, such as S/390 mainframes.

<http://www.alsa-project.org> – The official project page of the ALSA project, providing sound card drivers for the 2.5/6 kernels.

<http://kt.linuxcare.com> – A useful digest of all the major events on the kernel mailing list.

<http://www.x86-64.org> – Information about running Linux on 64-bit AMD processors.

<http://www.indrema.com> – Information about Indrema's open source games console.

→ Linux kernel ran to about 71K, supporting just very limited i386 hardware, whereas the latest 2.4.0test kernels run to approximately 21Mb. Despite this, however, it is not fair to accuse the kernel developers of bloating it unnecessarily.

As well as the number of features increasing, one of the biggest reasons for the increased size is due to the improved hardware support. The sheer number of hardware drivers now provided in the kernel is astonishing. This is in addition to the alternative hardware architectures that are now supported, going beyond the vanilla Intel/AMD i386 for which Linux was originally developed. The improved hardware support with the increased number of hardware architectures should lead to the improved scalability of Linux, and hopefully bring greater reliability with it. Kernel 2.4 really will be the kernel to use for cutting edge hardware, although it will support older hardware equally well.

New hardware architectures

With the 2.4 kernel, one of the aims was to improve the scalability of Linux, making better use of high-end machines – particularly in the server sector. The ability to support multiple processors is known as SMP (Symmetric Multi Processing). On x86 hardware, Linux has long had support for SMP, although this has been severely limited until fairly recently. Clearly, for enterprise machines, improved support for SMP is vital

In the 2.4 release, the number of supported processors on x86 hardware has been dramatically increased. Many commercial versions of Unix were more scalable than kernel 2.2 was, so this improved SMP support will help to

strengthen Linux's position as a reliable alternative to the commercial versions of UNIX, such as HP-UX, IRIX, and Solaris, and in fact, as a strong alternative to Windows NT/2000.

Kernel 2.4 brings three new processor architectures to the fold. Perhaps of most immediate significance to corporate users is support for the IBM S/390 mainframe. IBM itself played a major part in porting Linux to S/390, as it sees Linux as a strategic

Kernel 2.4 brings an interesting innovation in terms of web servers...

operating system. IBM has contributed the S/390 source code back to the community, and this has now been integrated into the regular 2.4 kernel sources (and to a lesser extent, into 2.2 as well).

Also of importance is support for Ia64 (aka Itanium), the name given to Intel's upcoming (but currently not commercially available) 64-bit processor that is likely to be released in 2001. Since the days of the 386 in the late 1980's, PC processors have been 32-bit only. Initially these chips will be mainly used in high-end machines, so it is of less importance to home and small business users. Eventually however, it is likely that 64-bit processors will become mainstream. Many large companies have contributed to the Ia64 port of Linux, including Intel, Hewlett Packard, and Red Hat.

The third processor architecture supported is the lesser know SuperH. Ironically enough, this is handheld Windows CE style hardware – eventually providing an alternative to

Window CE on this platform.

On the current 32-bit x86 hardware (e.g. Intel and AMD), the maximum amount of memory supported by the kernel has been increased to a massive 64Gb (up from 2Gb in the stock 2.2 kernel), taking the abilities of the kernel right up to the realms of enterprise level machines. For standard x86 processors, more clone processors are now explicitly recognised, and optimisation has been

provided for them. This is possible because portions of the kernel source code are written in assembler (although the majority is written in C), and this can be

optimised for particular processors. The kernel now fully supports the Crusoe (from Transmeta, the company which employs Linux's creator, Linus Torvalds), the AMD Athlon (aka K7), the Intel Pentium III/4 (with support for the new multimedia instructions), and various flavours of IDT's Winchip.

USB, Firewire, and other buses

The 2.4 kernel is the natural home for support for communication buses such as Universal Serial Bus (USB) and Firewire. The support (or lack of) for USB, and to an extent, Firewire, was a criticism regularly levelled at kernel 2.2, as many of the peripherals being produced now such as printers, scanners, modems, keyboards, and mice all use the USB interface.

USB controllers have been included on virtually all x86 motherboards since the days of the Pentium 133. Although USB was an Intel innovation, Apple have fully embraced the standard, with virtually all new Macintoshes supporting



IA64

Ia64 Linux aims to get Linux running on Intel's 64-bit processor

the bus. In fact, with the iMac, there is little room for expansion except with USB peripherals – both the keyboard and mouse are connected via USB. Additionally, Windows has been able to support USB since the release of Windows 98. Therefore, proper support for USB is seen as a key to Linux being accepted on the desktop. Support for USB was always planned for 2.4, although the kernel has taken longer to arrive than originally expected.

Firewire (sometimes known by its official name, IEEE1394, or by Sony's name, i.Link), is also an important innovation in the Linux kernel. USB is too slow (12Mbit/sec) for certain uses such as connecting devices like hard disks and digital video. Firewire is much faster (400Mbit/sec), and therefore may well be more important than USB in the near future.

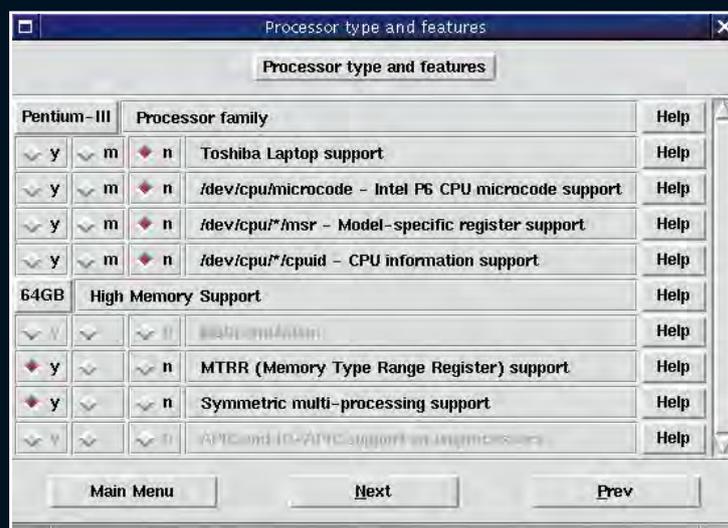
Internet Features

Kernel 2.4 brings an interesting innovation in terms of web servers – a specially designed web server known as kHTTPd is integrated into the kernel, as a kernel module. However kHTTPd doesn't aim to compete with normal HTTP daemons such as *Apache* and *Zeus*. kHTTPd was specially designed to

only support static content, in an attempt to dramatically increase throughput on web servers which support millions of hits per week. When a request for non-static content is received (for example a CGI script), kHTTPd will pass the request to a standard userspace web server, such as *Apache* or *Zeus*. kHTTPd is configured and controlled with the pseudo `/proc` filesystem.

The motivation behind kHTTPd was – partly – the results of comparison between Windows NT and Linux as a web server by Mindcraft. It later surfaced that Microsoft had probably paid Mindcraft to conduct these tests, and as a result it was perhaps no surprise that Windows NT came out as the winner. We won't go into detail here, but if you are interested, see <http://www.opensource.org/halloween/halloween6.html> for details of the supposedly rigged tests. The popularity of Linux as a world wide web server was further inspiration to produce kHTTPd, to make Linux an even better option in this vital sector.

Each new minor kernel version since 2.0 (including 2.2 and 2.4) have included a rewrite of the IP masquerading and firewalling tools. In



Linux become much more scalable, with the release of the 2.4 kernel

FEATURE | 2.4 kernel



Linux, support for firewalling and such measures are provided explicitly in the kernel, although the system is configured using the userspace tools `ipfwadm` (in 2.0), and `ipchains` (in 2.2). 2.4 is no different in that it contains a re-designed firewalling section, which is designed to be more powerful and secure than previously.

The new firewalling/masquerading methods are still configured userspace tools, in this case `iptables/netfilter`. One of the drawbacks of previous re-writes was that `ipchains` and `ipfwadm` used a different command line syntax for configuration purposes, making it less convenient to upgrade from a 2.0 kernel to a 2.2 kernel. However, the new `iptables` has been designed to have compatibility modules for both `ipchains` and `ipfwadm`, so there should be less short term difficulty when upgrading to 2.4 kernels.

IPv6 is the new implementation of the IP protocol (IPv4 being the currently used version), the protocol which is used to route traffic over the Internet or local networks. IP works by assigning a unique IP address/number in the form `nnn.nnn.nnn.nnn`, where 'nnn' is an integer value between 1 and 255. One of the drawbacks of IPv4 is that it only provided a finite number of unique numbers before running out. As every computer on the Internet

The standard 'make xconfig' screen of a late 2.4.0test kernel, showing the new hardware support

GLOSSARY

SMP – Symmetric Multi Processing

ALSA – Advance Linux Sound Architecture, a project to provide a high quality set of sound card drivers in Linux.

S/390 – A high end IBM mainframe computer.

x86 – A generic name given to a 384/486/586 or above, Intel compatible processor.

i386 – An abbreviation for the Intel/AMD 386 processor – often considered to be the lowest common denominator in hardware terms for Linux systems.

FEATURE | 2.4 kernel

→ must have a unique number, clearly there will be a limit to the number of computers which can be connected to the internet with IPv4. Although this is not yet a problem, it is likely to be a problem in the future. IPv6 was introduced to solve this problem, by providing many more possible combinations of addresses. Support for IPv6 was initially a little lacking, but in 2.4, it has been developed further, meaning that kernel 2.4 should be a solid choice for networks employing IPv6.

Filesystem support

A proper journailling file system is one of the features that Linux has missed previously, but there are now many competing standards, one of which is likely to become the standard in the near future. A journailling file system is different from a standard file system in the way that it operates, making it less prone to data loss. Ext2, the standard Linux data file

system, is reasonably robust, however it is still slightly prone to data loss in kernel crash and power loss situations. Clearly this is not acceptable in mission-critical situations, and in fact for home users, who may lose valuable data.

It was originally planned to have proper support for journailling file systems in the standard 2.4 kernel, although slippage has meant that this is highly unlikely to be the case in the first versions of the 2.4 kernel. The current two best known journailling file systems are ReiserFS, and ext3 (the natural progression from ext2), with both being reasonably robust. ReiserFS, however is possibly slightly more popular, with several distributions (based around the 2.2 kernel), such as SuSE and Mandrake providing it as an alternative to ext2. As ever, it is good to have competition, so which ever of ReiserFS or ext3 matures the first, is likely to become de facto standard soon.

In the 2.4 kernel, support for other (non-UNIX) file systems has been overhauled. Support for NTFS, the file system used by Windows NT/2000 has been improved, meaning that there is more chance of being able to read and write to partitions based on this file system. This is certainly of benefit to people who are dual booting Windows NT/2000 and Linux.

Conclusion

So, will the 2.4 kernel be worth upgrading to or using? Well, for almost everybody, the answer will almost certainly be "yes". With its mix of features for both home users and corporate users, the 2.4 kernel looks like it will be a major milestone in Linux scalability and usability. Even if you are not brave enough to upgrade a current distribution running a 2.2 kernel, distributions packaging the 2.4 kernel as standard are likely to be available very soon from all popular vendors.

LINUX: THE NEXT GENERATION

The release of the 2.4 kernel will be a very big event for both the Linux community, and to an extent, the entire free software community.

The kernel has improved immeasurably since early versions, although the pace of innovations in other areas of the IT industry has meant that development has had to continue at a quick pace.

Most people will be happy with it, although it leads onto the interesting question. What is likely in the next generation of Linux kernels (the 2.5 development kernel, and eventually the 2.6 stable kernels)?

Multimedia

Increased or changed hardware support is likely to be a large part of the next kernel release (the 2.5 development

kernel, and eventually 2.6), making Linux more suitable as a desktop operating system. One of the most visible projects at the moment is the ALSA (Advance Linux Sound

One idea is to overhaul the method of recompiling the Kernel to make it more modular

Architecture) sound card drivers. It has already been decided that these will be integrated into the kernel, in place of the current standard OSS/free (and to an extent, the commercial OSS as well. OSS/free is limited in both the number

of cards that it supports, and the sound reproduction qualities of the drivers. Therefore, ALSA is designed to provide answers to both these problems. ALSA is also designed to be more modular.

The development of ALSA will provide several new APIs (Application Program Interfaces) which will make it easier to create programs that support sound.

Linux is also likely to need to develop to support new mass storage media such as DVD-RAM. This is currently expensive, and out of reach of the mainstream buyer. But by the next release, such storage mediums are likely to be common.

Another interesting prospect is the outcome of projects such as Indrema (see issue 5, which includes an interview with Indrema's co-founder) – a company which aims to create a games console based around Linux. If Indrema make any modifications to the kernel source code to suit their need, these must be distributed back to the community, under the terms of the GNU General Public License. Modifications are likely to improve gaming performance, so if these modifications made by Indrema are integrated back into the main kernel, perhaps by the time of 2.5/6, the kernel is likely to be better suited to games playing. A similar story exists with Nokia's Media Terminal (covered in our story on IBC 2000, News, issue 7), which is based around a similarly open core.



The ALSA project will be one of the major contributors to drivers in the 2.5/6 kernel

to bring out their own version of a 64-bit processor (codenamed Hammer) in the future. The 64-bit Intel and AMD chips are likely to be quite different, with less compatibility between Ia64 and Hammer than with Intel's and AMD's 32-bit chips. Therefore more work will be done on porting Linux to Hammer. AMD have already started the ball rolling in this respect, with the production of a web site housing tools such as a developers toolkit, consisting of preliminary versions of essential Linux tools such as GCC (GNU C Compiler), and binutils, with some help from SuSE. An interesting press release on the subject can be seen at <http://www.amd.com/news/prodpr/20163.html>.

Networking

NFS (Network File System) is the standard method for sharing files over UNIX connected networks, and of course Linux supports this well. Version four of NFS will eventually be supported by Linux, likely to be around the time of kernel 2.5/6. The 2.5/6 kernel will almost certainly have continued good support for SMB (Server Message Block), the protocol used by Windows networks to share files and printers.

Hardware and scalability

As you have already seen in this article, the scalability of kernel 2.4, on x86 hardware, has increased to a large extent. This is both due to the increased SMP configuration, and the greater amount of RAM supported. Ia64 (Itanium) will be Intel's version of a 64-bit x86 processor, however AMD are also intending

maintainable for the developers. Suggestions such as these crop up fairly regularly on the main kernel mailing list – take a look at the digested archives available at <http://kt.linuxcare.com> for more information.

Another possible idea is to overhaul the method of recompiling the kernel. Currently this is not a particularly easy step, with several commands needed, so a better method would be welcome. Details appear patchy about this project at the moment, but if you are interested, see <ftp://ftp.kernel.org/pub/linux/kernel/projects/kbuild> for more information.

Fragmentation?

There are always some people voicing concerns that Linux will fragment, and soon we'll all be running incompatible software on different Linux spinoffs. The headline here is that some fragmentation is good!

The ability to use Linux successfully for embedded solutions depends on some fragmentation – and it can only benefit Linux as a whole to be used in a variety of exciting and new ways.

The pace of change

Some time ago, Linus Torvalds expressed an interest in increasing the rate of kernel releases. Unfortunately, Kernel 2.4 will be released later than it was intended to be, so this dream has not yet been realised. Kernel 2.4 will go a long way to solving many scalability issues, however there will always be other issues that need addressing, leading to continued kernel development. Who knows – by the time of the 2.5/6 kernel, Linux could well be a standard, and perhaps be even more popular than the commercial versions of UNIX. With the continued support of industry heavyweights like IBM, and even "competitors" like Sun, this seems pretty likely. 



Linux kernel 2.4 has improved support for several IBM mainframes

Apache has had its own way for a long time, and **Dave Coulson** thinks it's time the competition put itself about a bit.

Roll over

If someone asked you to name all of the open source web servers, how many would you come up with? There is, of course, *Apache*, the original NCSA http daemon, and maybe a couple of the simple servers such as *boa* or even the *khttpd* from Linux 2.4. Very rarely will someone name a web server which has been in the open source community for many years and is a major contender to the likes of *Apache* and Zeus' server for ASPs. This server is, of course, *Roxen*. Technically, the web server is called *Roxen WebServer* which is a bit of a mouthful, so people usually just call it *Roxen*. Originally it was part of a larger group of applications known as the *Roxen Platform*, which included the web server, at that time called *Roxen Challenger*, along with a high power log analyser and a few other things which people didn't really care about. These disparate apps have since merged into a single application suite which includes the web server, as part of a commercial web development package. Fortunately for us, the most useful part of the suite is freely available.

Roxen is built around a programming language called Pike, which is loosely based on C, except for the fact that it is completely object orientated. All of *Roxen's* modules, are written using Pike, which makes creating your own *Roxen* modules a lot easier than *Apache* ones, simply because Pike is a lot simpler than C to code in. However, the clever chaps who built *Roxen* realised that people don't want to learn Pike to have fun and interactive features in web pages, so they created their own XML style language, which you can simply embed into the HTML, called RXML. RXML is so fantastically simple that even someone who've never programmed a line of anything in their life can manage to produce something which would require a lot more code and a lot more understanding in, say, PERL or PHP. Add to this the fact that *Roxen* modules can supply their own RXML tags, and you end up with a simple, yet extensible, web programming language. Of course, if you want to do really complicated things, such as use databases, you will need to know how to use SQL and such like; *Roxen* can't help with that, but when it comes to doing the actual query, everything is nice and straight forward.

As RXML is simply a high-level layer on top of Pike, you can include portions of Pike code in web pages, which can access the same features as the RXML script, as well as any RXML variables or information within the page. This gives increased flexibility when you want to include a small portion of complex code on a single page, but don't fancy rewriting your few hundred lines of RXML into a couple of thousand lines of Pike.

We don't want to run before we can walk, so getting *Roxen* to serve static HTML is what we will be looking at first.

Getting and installing Roxen WebServer

Downloading the latest version of the WebServer is easy (the latest available is included on our CD, but it is updated often). Just hop over to <http://download.roxen.com/> and select WebServer 2.1. Whether or not you are

able to download the 128bit crypto version of the server depends upon your location, although people within the EU should not have any problems. The whole package is around 9.5Mb, so you may want to go and make a coffee while it downloads. Or get it from our coverdisc and save your phone bill.

At a first glance, you may say "This ends in .sh; what have you made me download?". The shell script contains all of the packages needed to install *Roxen*, so you don't have to go out and download individual parts of the server, such as Pike, separately.

Installation is a no-brainer; Just type `./roxen-2.1.185-lrh6.2.sh` and you can pretty much keep to the defaults. After it has installed, it will ask you for a password for the 'administrator' user. Make sure you remember this password, as it's near to impossible to recover the server if you forget it or type it in wrong.

When it has finished installation and basic configuration, it will return a URL, such as `http://localhost:11466/` which is the URL to the web administration tool. After logging in with the username and password you supplied during installation, you end up at the pretty admin tool. Everything you need to set up, maintain and generally administrate the web server is available from this web based tool, so there is none of the messing around with configuration files as with most servers, but that's not always a good thing, particularly if your trying to configure a server over a 56k modem or something equally nasty.

From here, we need to set up a 'site', which corresponds to a virtual host. Of course, before we can start to set up a site, we need to have the web server listening on a port other than the exclusive admin port. Usually, web servers listen on port 80 for http and port 443 for http over SSL. As we already have *Apache* listening on port 80, we're going to run *Roxen* on port 8080, which, although usually reserved for a web proxy, generally doesn't interfere with people who have restricted Internet access.

To create a virtual host, just click the 'Sites' tab at the top, then 'Create new site'. Give it a name, which doesn't need to have anything to do with the hostnames it is handled by, then click 'Create new site with template' and pick 'Standard Site'.

The first of the settings is URL. In here, we put the hostname we want to handle this, in this case 'taitiu:8080', the latter portion being the port number.

You can, if you choose, have it listen on more than one IP and more than one port, but it's unlikely that many people will use this feature. The only other field we're interested in at the moment is the 'Search Path', which is the same as the DocumentRoot in *Apache*. We're just going to use '/home/httpd/htdocs'. You can, if you want, have more than one Mount Point and Search Path, a little like an Alias entry in *Apache*.

Once we have created our server, we can just click on the link at the top and it will open a new browser window containing the data from the requested server. While all we've done was click on stuff – and most people familiar with *Apache* could have probably done it in half the time – someone who may not have a clue about setting up web →



the world

FEATURE | Roxen

→ servers has now got one running on their machine. Just plonk an index.html into the directory which you chose as the Search Path, and away you go. Simple, wasn't it?

At this point you might want to look at the list of modules to see what is on offer, as many of them are quite obscure and aren't really comparable to similar systems. There are a lot of filesystem, authentication and graphics modules, all of which you can load on demand and configure as you choose without restarting the web server.

Roxen Features

Now we have our own server setup, we can start to add specific modules to it's configuration, so we can use their features in RXML and Pike scripts.

First of all, we want the SQL module. Just add the 'SQL Tags' module and select the correct default SQL server. For example, if you have mysql running locally and want to access it with the username 'dbuser' and the password 'mysql' you would do set the default server to 'mysql://dbuser:mysql@localhost/defaultdb'. Once you enter this information, it will take you to the Information page which will show the status of the connection and the reason for failure, if any. Quite conveniently, the module also lists all RXML tags which have become available by the inclusion of

the SQL Tags Module in your site, so there is none of that hunting around trying to figure out what the modules can do.

We want to check to make sure that queries work right, so we need to start writing our first RXML script. As mentioned earlier, RXML has exactly the same format as XML, so it's not particularly difficult to adapt existing pages to show some dynamic content. As default,

Roxen will parse .html, .htm and .rxml files for RXML content, so we're just going to use index.html to test RXML scripts. As you have probably already noticed from the tag list, there are four different tags which do SQL functions, so which one do we pick? Many RXML tags have built-in help functions which can be printed on the web page, by typing `<tagname help/>`. For those not familiar with XML notation: if a tag is just a single tag and not a container, then you put a / before the final `>`. We are just going to make it spit out a list of all available databases, and there are two ways of doing this.

The first is simply a method of printing SQL data to a page; `sqltable`. We would do

```
<sqltable query="SHOW DATABASES"/>
```

and reload the page. Simple? The second option is a lot more flexible, as it allows us to insert the returned row data into whatever HTML formatting we like, including sending the output from one query into other RXML functions. This is known as `sqloutput` and has a slightly more complex usage than `sqltable`.

```
<sqloutput query="SHOW
DATABASES">
#Database#
</sqloutput>
```

As the SQL server returns the rows, the RXML parser will loop the code within the `<sqloutput>` container. It stores the returned field names as RXML variables, which we can print

The features most people choose Roxen for are the fancy graphic tools

to the page by putting them within `#` symbols. From these simple examples, it's relatively easy to expand them into, say, a news section on a web site.

Not all SQL queries need output, such as UPDATE or INSERT queries, so there is a `<sqlquery>` tag which simply executes a query and nothing more. `<sqlquery help/>` has more information if you want to use that function.

Graphic functions

As much fun as SQL functions are (hah!), the features most people choose Roxen for are the fancy graphic tools and functions. Standard sites have the 'Gtext' module installed as default, although the 'Business Graphics' and 'Graphical Counter' modules will have to be added via the admin interface. There is also a nifty little 'Atlas' module which can be used to generate world maps on demand, although I can't imagine many people finding a productive use for it.

Gtext is great fun, not to mention is rather inefficient use of CPU time. Making a title or a headline is dead easy, so there is no longer the excuse of lacking artistic talent. As with other RXML tags, `<gtext help/>` provides a rather complete list of all Gtext options, so even the most incapable artist can produce something which looks almost decent (but don't come blaming us or the Roxen developers if it doesn't).

The most basic gtext tag is

```
<gtext>Any old text</gtext>
```

which will produce a fairly plain style text output. Notice that if you have not declared a background colour, either in a `<body>` tag or within any table tags, it will assume that the background colour is white, even though it is grey in Netscape. If you want it to be black on grey, make sure you have it set the background to grey in the HTML, or it will look very nasty.

We obviously can't go over all of the Gtext options in any



A web front-end makes administering your site with Roxen an easy prospect.

great detail, but the documentation from the module is enough for anyone to understand how to use it.

As much fun as making pretty looking text can be, it's not usually that much use. The 'Business Graphics' module, or the 'Graph' module as it's often known, is a truly excellent example of what you can make *Roxen* do. Coupled with the SQL modules, you can't really do much better than *Roxen* if you want to create graphical output from SQL data or user input. Without the time to setup some database tables and bits, we're just going to generate a simple graph using some static data.

We're going to plot a line graph containing some example data, then put a nice table underneath showing what each line refers to.

```
<diagram type='line' width='500' height='500'
name='Example Data' horgrid=""
tonedbox='lightgreen,lightblue,white'>
<data form='row' separator='.' xnamesvert=""
xnames="">
Jan. Feb. Mar. Apr. May. Jun
25,000.26,500.27,000.23,000.25,000.21,000
12,000.15,000.14,400.18,000.21,000.25,000
18,000.21,000.28,000.32,000.28,000.24,000
28,000.32,000.34,000.33,000.35,000.27,000
</data>
<legend separator=', '>PCPlus, Linux Format, .net,
PCFormat</legend>
<xaxis quantity='Month'/>
<yaxis quantity='Some arbitrary value'/>
</diagram>
```

It all looks rather complicated, doesn't it? Sadly `<diagram help/>` doesn't work, probably due to the extortionate amount of options you can use, so you'll have to resort to the documentation on roxen.com if you get stuck. Suffice to say, the business graphics module supports almost all simple graph types, as well as a couple of fancy 3D ones, so there will be something there which displays your data in a nice fashion whatever it should be.

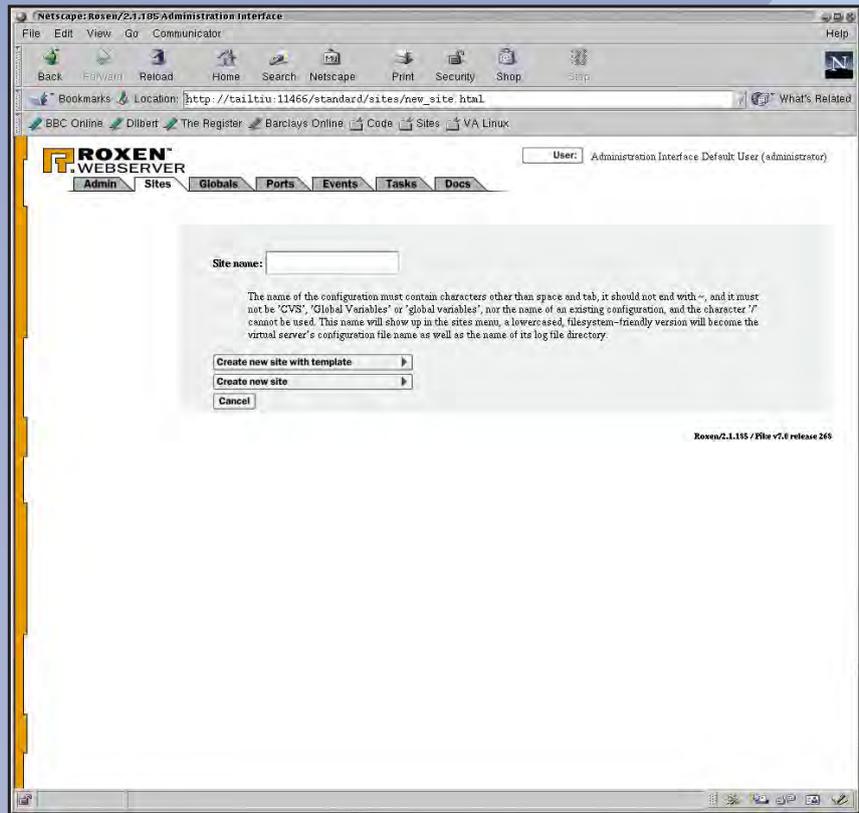
Real Coding

Making pretty pictures and doing SQL queries is fine up to a point, but what happens when you want to do something complicated in the RXML code?

Don't worry, RXML isn't completely stupid. You can do loops, store values in variables and insert them into other RXML functions, define your own RXML tags.

First of all, we need to understand how variables work in RXML. Each page contains various scopes of variables. These include **&form;** which contains variable values from forms, **&var;** for general variable storage, **&client;** which contains client information, such as which browser they are using and **&cookie;** which contains cookie values. Setting a variable is done using the `<set>` tag in the following way

```
<set variable="var.foo" value="something"/>
```



As this is in the same format as XML, the variables are defined as XML entities, rather than forcing us to use another tag to extract variable data. You may also have come across these in HTML as **"**, **&**, **<** and such like. So, to print the value of the variable 'foo' in the scope 'var' we do

```
&var.foo;
```

Simple really, isn't it? We can of course include that into another RXML tag, so

```
<gtext>&var.foo;</gtext>
```

is perfectly valid usage.

Conditional statements are almost as easy, although there are two different types. If we want to compare two variables to see if they're the same, we use the `<if variable="...">` tag.

```
<if variable="var.foo = something">
True
</if>
<else>
False
</else>
```

Using our previously set variable, it would print 'True' within the HTML. We can, if we want, have it execute more RXML functions, rather than just print something out. Note that `variable=""` within `<if>` compares a variable to a string, so if we wanted to compare two variables we would have to use something a little different

This is the first step on the path to creating your first Roxen website.

FEATURE | Roxen

```
<if variable="var.foo =
&var.test;">
True
</if>
```

Here, we've replaced the static string 'something' with the contents of the 'test' variable in the 'var' scope. As the entity is included before the RXML tag is parsed out, we can include as many of these as we want

```
<if
variable="var.foo=&var.test;&var.moo;">
```

and such like.

If we want to do a simple pattern match on the variable we have to approach it slightly differently. Rather than comparing the variable to a string, we want to compare two strings, which requires us to insert the entity, rather than the variable name when doing a match.

```
<if match="&var.foo;= *thing">
True
</if>
```

Roxen only does simple pattern matching, so we're stuck with using * for two or more characters, or ? for one character.

We can also expand on the `<if>` statements with `<elseif>`, so we can produce proper conditional functions from a few simple RXML tags.

We can also do loops in RXML, using the `<for>` tag. It's really easy to use



Graphic text is a doddle with GText tags

```
<for from=1 to=51 step=2 variable=foo>
```

Here, we've just use 'foo' as the variable, rather than 'var.foo'. We've decided to create a new variable scope called 'moo' to demonstrate the use of localised variables. We can define a variable scope using

```
<scope name="moo">
...
</scope>
```

and anything within the container will be within the 'moo' variable scope. Anything outside of this scope won't know about it's variables, but this scope will know about variable scopes within it's parent. An example is probably in order to demonstrate that.

```
<set variable="var.foo" value="something"/>
<scope name=moo>
<for from=1 to=51 step=2 variable=foo>
&var.foo; <insert variable=foo/><br>
</for>
</scope>
```

This will print out 'something X', where X is every odd number between 1 and 51 inclusive. Note that we need to use `<insert>` rather than the entity style when something is defined within a non-global scope. While the 'moo' scope can access variables from the 'var' scope, once you break out of the 'moo' scope, you can not access it's variables, so adding `&moo.foo` to the end of that script results in a 'RXML parse error: Unknown scope "moo"'. However, even if we reenter the scope 'moo'

```
<scope name="moo">
<insert variable=foo/>
</scope>
```

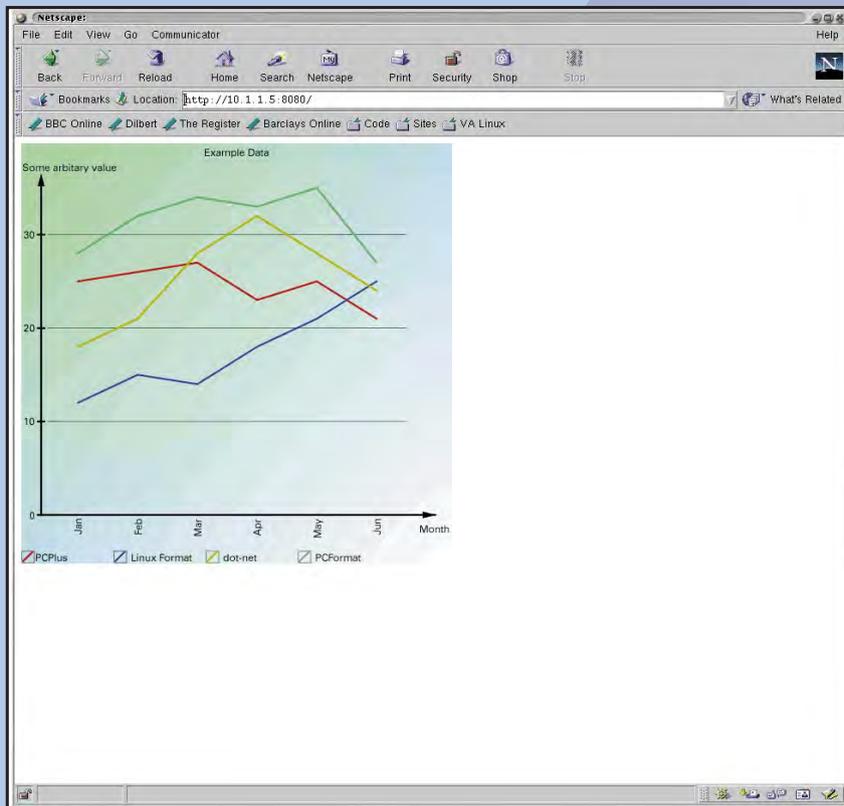
it will not return anything, as the 'moo' scope was previously removed by the `</scope>` further up in the script.

There are many more RXML functions, although they all stick to the same syntax and rules, and given the level of the documentation available, both interactively and at roxen.com, there is always somewhere to turn for help.

The price

Flashy text rendering functions and easy to use XML style

Graphs are slightly more complex, but are still simple to create.



parsable languages are nice, but they do unfortunately come at a cost. As you can probably imagine, generating little pictures of the text is not great for either the server, or the poor guy (or girl) paying for your bandwidth, so you probably won't want to build a page entirely out of `<gtext>` tags. *Roxen* does contain excellent automatic caching code for all of its image functions (good job too, it's got an awful lot of them), but if you have dynamic data, such as that from a database, the caching will only help up to a point.

Parsing RXML isn't much fun either if you're a server, so unless you really need RXML, make sure that the server isn't going to start parsing all of your flat HTML pages. As default, *Roxen* will parse .htm and .html files for *Roxen* tags, which, in our opinion, is a big mistake. .html should be reserved for static content, and use a different tag, such as .xml, for those which require parsing. It's comparable to shoving all of your HTML scripts through the SSI parser in *Apache*, although obviously to a rather larger degree.

Apache or Roxen?

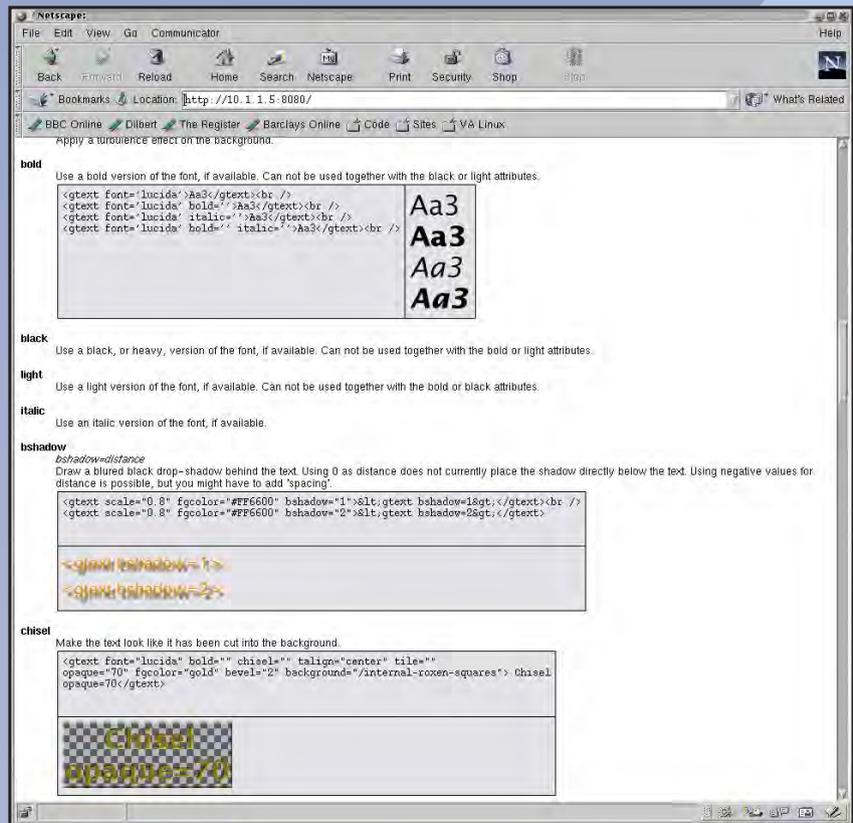
At the end of the day, it boils down to a simple question; "Should I use *Roxen*?" For most people, who just use a web server to power the local copy of their web site, or just to mess around and learn things, then *Roxen* is a good choice. It has lots of little bits and bobs which can keep the inquisitive mind occupied for hours, although you have to bare in mind that your ISP probably won't be running *Roxen*, so don't get carried away with its additional features.

For a high-end production web site, *Roxen* offers comparable performance to that of *Apache* for static files, although neither match the performance of *Zeus*. *Roxen*, however, does not have many of the highly specific modules which *Apache* does, such as `mod_perl`, which will inhibit the ease of transfer from *Apache* to *Roxen* for many people. However a major plus is that PHP is able to be compiled as a Pike module, so if you use PHP for your site, you just build PHP as you would for *Apache*, and plonk it into the site configuration and off you go.

In the real world

We find that *Roxen* is ideal as a back-end server, with the remote requests going via *Apache*, along with all of the specific modules we need to drive various sites, then use internal fetches to pull dynamic content from *Roxen*. So, for the sake of a little bit of Perl code which does the internal fetches, we get the best of both worlds. To avoid wasting CPU time on *Roxen*'s generation of images, the Perl script stores all of the fetched data on disk, so even if *Roxen* decides that it should really regenerate the whole thing again, it doesn't get a chance to as it is not called again until the cached content times out. Obviously such as setup is a bit of a dirty hack, but at least everything works as it should.

It's not just the features of the web server which can cause problems; it's the code which people try to make it run which is a cause for concern. Being such a simple and easy to write language, anyone can come along and write a piece of RXML script which pulls all sorts of things from databases and plasters them around the page. That's fine, as only you can



update the content, so it will only change when you force it to, but what if you had a simple form which took someone's name, and generated a piece of graphical text from it? It, at first, seems pretty innocent; the sort of thing someone might put on their site as a bit of a laugh, but it's enough to form the basis of a really simple denial of service attack which would take down every site hosted on that particular *Roxen* server installation. I'm not suggesting that you couldn't do the same with *Apache* or *Zeus*, but if you aren't aware of the risks, you may be more vulnerable

Conclusion

You're obviously not going to get a "Use *Roxen*" or "Don't use *Roxen*" statement from this article, so it's really up to you to evaluate your needs and decide if *Roxen* is right for you. Though it is a less popular server than *Apache*, *Roxen* users have built up an extensive online community, so it would be in your interests to join some of the mailing lists. You may not understand everything that is going on at first, as people try to figure out how to hack a Pike module to do something strange, but there will – more often than not – be someone there who can fix any problems you're encountering with *Roxen*, or offer advice as to its usage or configuration. Usually with *Apache*, you resort to asking people you know, or heading to a specific module mailing list. It's far more pleasant, particularly for a beginner who just wants a web server to serve the odd page here and there, rather than conquer the world, to have a single place where they can get advice from people in the know and pass that on to others.

If you get really lucky, one of your suggestions might end up in the next release! **LXF**

Roxen has an ingenious help system: just include <tagname help/> in your code and voila, all the options.

Web resources

The first place to go for Roxen news and goodies of www.roxen.com.

http://community.roxen.com features articles, news, forums and developer tools.

www.devshed.com/server-side/roxen has a number of Roxen features and a fairly useful intro to the software.

http://pike.roxen.com is the place to go for all things Pike including modules, source code and news.



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

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. 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! It's surprising how often we are left wondering who exactly "bounding_mantis@hotmail.com" really is. We'll start making up silly names soon, if you aren't careful!

Nick Veitch
Editor

Long names

I have just bought LFX08 and once again have enjoyed the content so far. I have just one minor complaint/problem/request.

I do not have a CD-ROM in my Linux box, only in my Windows PC, so I have to transfer the files from your CD via the network to my Linux system. Both Samba and ftp can cope with long file names but the files on your CD are always truncated to the dos file name format (8.3).

If you can't write the CD with long file names could you include a translation of which 8.3 names map to the correct names. As an example I have two files in the root of this months CD called Smoothwa.000 and Smoothwa.51f, both are 1.440K one must be the SmoothWall boot disk and one must be the module disk, with trial and error I could figure it out but it would be easier if you could provide a translation so the instructions included would make a lot of sense.

Apart from that, keep up the good work.

Dave Moor, via email.

Sorry about that. Actually, although we (naturally I think) construct the disks primarily for use on Linux, we do normally include joliet extensions for full filenames under Windows, and I'm sorry it was missed off that month's cover disc.

Back issues

I'm interested in getting hold of some back issues of *Linux Format*. At the moment, I only have issue 6, and issue 8, but would like to complete my collection. How much do back issues cost,

and how do I go about getting hold of them?

Brian Heys, via email

For you, and for all the other people who write in about back issues, you'll now find a back issues page near the back of the mag (page 97 this time). Unfortunately, they do sell out occasionally.

Identify yourselves

We are trying to find the ISSN number for *Linux Format* magazine. We can't find it within the magazine and we were wondering if you could tell us please.

Tennille Patchell, Canberra, Australia

Hi there, The ISSN for you, or anyone else wishing to place orders at bookshops is 1470-4230. A tip for you here is that on most UK magazines, the ISSN number is the 8 digits following the "977" at the start of the barcode number.

Modem madness

In view of the fact that Linux makes an ideal Internet box it has always seemed strange to me that people have problems with modems. The problem seems to arise from the fact

that most so called hardware modems are only really suitable for use with Windows, making them work with Linux is fraught with headaches.

I have tried all implementations of Linux and never once had a modem problem, but then I always use a Dynalink 56K/V90 ISA internal modem with com port configured with links. The question then arises where are these modems obtainable?

Well a little known fact is that Dynalink (012 52 727711) can supply these modems to retailers BUT only up to February 2001. After this date production will stop, apparently due to the forthcoming unavailability of the Rockwell chipset and mother boards with ISA slots – what WILL we all do then?

The model number to look for is either MX56V-ISA or V1456VQH-ISA, both modems are the same it's just that Dynalink have recently changed the model number.

I have three of these modems running with both Linux and Windows machines and intend to stock up with a further two for future developments – after all they can be used with all operating systems, even Windows! My usual supplier is Ronal Computers in Southport (01704 507 808)

I have spread the word among my Linux friends and thought your readers might be interested in knowing how their modem problems can so easily be fixed. But beware, it's a case of buy now or miss out!

Arthur Binning, via email

Thanks for the tip, I'm sure this may help some readers. The

And for once, nobody wrote in and complained about the stickiness of the glue...



Some people like using KDE, some people don't, but is it destined to become the standard desktop for Linux?

problem is that most modern internal modems are what are collectively known as "winmodems". Some older models of internal modems, particularly ISA models, are configured to appear as straightforward serial devices, so they will work fine under Linux, which seems to be the case with the particular model you are using. There is plenty of information on the Internet for people wondering whether or not their internal software modem will work with Linux at www.linmodems.org

Long Rant

In reply to Rob Johnson's (and other) email(s) in the Christmas 2000 issue. I use *Window Maker*, not *GNOME* or *KDE*. Because I like it. That's not to say that *GNOME* and *KDE* are rubbish (just bloated), I just prefer using *Window Maker*.

Yes M\$ Windoze has a nice GUI, but that doesn't make them innovators in GUI OS design. Far from it, most of the ideas within the Windoze environment are borrowed from X or even the Mac OS.

What makes Windoze so accessible then? Easy, it's what people are used to. Even if you are a newbie to Windoze (and lets face it, we have all been there) and you have a question, nine times out of ten a mate will know the answer OR knows someone who knows the answer. So really it's about home users swapping to Linux and getting used to a new environment. Quite obviously this is happening, if a little bit slowly for some people.

So, lets be realistic about this, when Linux becomes the games platform of choice, how many people are gonna want us computer types to install Linux on their machines? Everyone, bar your Ma and Pa. And will it really matter to them that they are using *kmail* and not (spit) *Outlook Express*, *Abiword* instead of *M\$Word* or *gnnumeric* instead of *Excel*. No.



Because they will be running *Quake III Arena* about 20% faster with 100% fewer crashes.

One thing that will matter to these people is a familiar environment, hence there is nothing wrong with

removing Windows from your computer, because I have recently done just that. I had not, after all, booted Windows in something like six months, and was not likely to, either – it involves a blown line transformer in my original monitor, an ancient 14 inch SVGA replacement and the fact that I really couldn't be bothered to figure out how to boot in safe mode in Windows. That and the fact that I don't know the vsync & hsync for this thing – it's so old that the manual is practically unreadable. Generic SVGA seems to work OK under X, though... And I had these shiny new

grounds that "file already exists", which the Uni boxes do regularly. You might be able to solve it by deleting the Internet cache, or logging out the back in again, or trying again later. Or probably not.

Meanwhile Mandrake 7.2 has all sorts of neat stuff to play with – including CUPS, which means the my Epson 740 is now (IMHO) better supported than it was under Windows and is certainly more configurable :-).

Now if only I had a 3DFX card...
[Sam, via email](#)

Thanks for sharing your experiences. As some people have pointed out, their dependency on Windows usually lies in the fact that there are some, usually business-oriented, applications they simply can't run under Linux, and there are no viable Linux alternatives. We'll be looking at these problems and a number of ways to get around them in a future issue.

Newbie stuff

First of all, Great Mag! But! You most likely get loads of emails and letters asking for help for the NEWBIES, well, this is another but with a difference. Having toyed with the idea of Linux for years, along came your magazine with a concise explanation of how to format and partition the hard drive and install Linux. Having read the article several times over, and now with a hard drive that has just been formatted (to clear the junk that Windows had gathered over the years), I took the plunge and installed Definite Linux. I was hoping for an account by account instruction on this but sadly it didn't appear, some of the simple things were not explained at all (like how to setup a new user) and many hours were spent trying to find the relevant information.

Now, I am not getting at LXF personally, because this is a problem that is in evidence everywhere, be it IT, or even

Preaching to the converted? Maybe, but it has prompted a lot of letters about the merits of dual booting systems.

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

designing a GUI to look and feel like Windoze (obviously without the half hour restart feature). I am glad for *KDE*, its an entry level window manager. When I introduce Linux into the administrative pool at work, they WILL be running *KDE*. *GNOME* comes in a close second. Why? Because it has the same feel as Windoze, but isn't as friendly as *KDE*.

So this is for my friend Rob, tell those *KDE* haters to check out *Fvwm 95*, and then go get a life. If they really wanna rubbish *KDE*, rubbish the libraries that it uses (QT), I believe that they are not true open source, am I right?

PS. Can we please end this thread here, I really don't want to read another article about how Linux is going to replace Windoze. Its going to and that's the end of it.
[erbi obi, via email](#)

Somehow, I doubt that will be the end of it...

Dumping Windows

I was interested to see the article in the current (December) ish on

Mandrake 7.2 CDs, so I thought might as well reorganise my partitions & switch to Reiserfs.

As it is, I have plenty of opportunities to compare the competition, seeing as how all the lab computers at University are (Ack! Pfft!) Windows NT. Working with *IE5.5* has even made me appreciate Netscape anew – it may disappear without warning, but it NEVER plays the trick of refusing to 'save target as' on the



LETTERS | your views

cars. Those who know, forget how difficult it is for those who don't, or for those who are trying to learn! The tutorials, so far have been interesting and I have managed to get about 30-50% info out of them but I believe what you need is a NEWBIE on your staff of writers who can put things across as they learn them and in the way they teach them to others. Another important but overlooked area is 'Where' to find things, i.e. /var, /usr, /dev, etc. What's in them? If I was to load a new program, where does it go? How do I put a shortcut on my desktop, or how do I get a link in the Start menu? You know how, I don't! A possible stupid example; do you always have to log out every time you want to go to "root" when in KDE or GNOME? Or is there a method of accessing "root"? And I don't mean through an Xterm! Why when you change settings in root, does it not automatically change in the other users area too? For me, these are all basic fundamentals that I believe all newcomers need to know, to help them work with Linux. Imagine my joy the first time I got "chmod" to work, so I could gain access to Kppp!

I tried Definite, linuxconf was great but I got deeply in trouble and had to reload. Then Storm, good but it didn't find much of what was on my computer. Then came Mandrake 7.1, wonderful but again didn't find all my internals, but SuSe 7.0... well what can I say. WONDERFUL! SuSe found just about everything, except for my USB smartcard reader. It even found my LS120, though I cannot! Keep up the good work but a bit more detail would be very appreciated.

J.Faulk, via email

Well, we do have a newbie on staff now, in the form of our Production Editor, Andy Channelle. He may not be a newbie for too long though, as he has to read every word of the magazine each issue!

I appreciate your problems, and I'd like to think that the magazine is addressing some of these concerns. We can't cover everything for newbies each issue, because the magazine

would probably have to be 30 times bigger, and nobody would read it more than once. Most of the questions you've raised in your letter have actually been covered in tutorials and in the answers section of the magazine. But we will try harder to explain things more clearly in future.

Applix, BSD and more

I have been using Linux for a year and enjoy reading LXF each month. Generally I find your magazine informative and helpful, but I was a little annoyed by a curious omission in LXF 09.

Why did you leave out *Applixware* from your round up of Linux word processors? This seemed to me a strange decision as it was, perhaps, the first to be ported (from UNIX!) across to Linux. OK, it isn't free and it isn't open-source but I have found it both faster and more stable than *StarOffice* or *Wordperfect 9*, and surely it deserved a mention at least! (Did you really do your tests on a machine with only 32Mb of ram?)

Also could you please do an article on BSD and UNIX generally and perhaps also examine the peculiar relationship that currently exists between UNIX and GNU/Linux and the differences between the two systems. After all, Linux is, in effect, a type of UNIX but without the big price tag and the special hardware, and we readers don't often get to read about UNIX in the usual computer magazines.

PS. I also believe that the reason for Tux the penguin (by Larry Ewing) being adopted as the Linux mascot was that it was one of the first quality pieces of artwork created on a Linux box using *Gimp*.

Andrew Richardson, London.

The omission of Applixware was a bit of an oversight, although we were concentrating mostly on the open source solutions. Well spotted! The roundup tests were performed on a machine with 320Mb of RAM.

BSD is an interesting subject. We are actually planning a look at Linux's close cousins in 2001, so keep reading, and thanks for the feedback.

There's life in the old dog yet. Don't throw out your old hardware, put it to use.

Mucking about

One reason I'm interested in alternative OSs for my A1200 is that since I got online one of my favourite hobbies has been MUCKing. MUCKs are a bit like MUDs, only they're not. (The acronym allegedly stands for Multi User Character Kingdom). I have been thinking a bit about setting up a MUCK at home, mostly for experimenting and learning without having to pay to be online. Actually I have already downloaded the files I seem to need, but I thought that other readers may find it interesting, captivating, maybe even an addictive past time. So my wish for next issue is: Fuzzball.

Ragnar Fyri, via email

Your wish, our command – we'll bung all the relevant stuff on the next available CD!

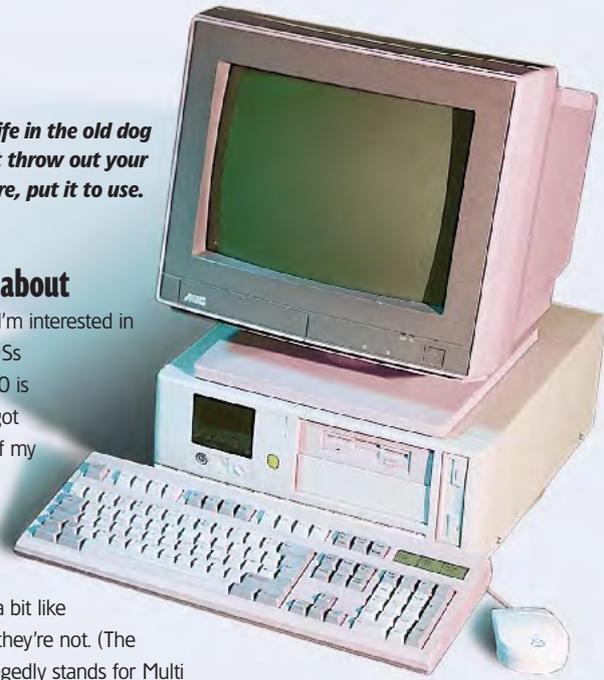
Old hardware

Thank you for such an interesting article (Making use of old hardware – Issue 9). I hope that folk make good use of it.

I have been using similar legacy hardware under Linux for some time. Amongst other roles, I manage IT for a small electronics company.

Our main server is a Novell *Netware* box (i486DX100). This is a bit of kit which has remained unchanged since 1993 (except for adding extra disk space) and performs with a level of reliability matching that of Linux.

However, as Novell 3.12 only supports 8+3 filenames and many of the desktops run Windows (in one incarnation or another), it has become necessary to support long file names and we are switching from IPX to IP networking.



I had available several "old" i486/[25,33] boxes and used these, running Linux, as file servers, intranet servers and even as the Internet gateway. I use Slackware 4.0 and do not try to keep at the bleeding edge – currently a 2.0.35 kernel is used with little or no problems. Similarly, older Sun kit running RedHat pops up in a similar role – even an IPC as a dhcp server.

One critical aspect of all these machines is RAM – I use 16Mb (24Mb on the Sun machine) as a minimum. (Obviously the HDD's are appropriate to the job.) Critical too is the kernel build and the load on the system, no unwanted functions like X, ppp, innd or the like either in the kernel or running.

The result is a system with high reliability at minimal cost. Performance is more than adequate to keep up with our 10baseT network and PII desktops. Uptime – limited by the local mains supply.

Geoff Blake, Chelmsford

There are certainly a wide range of tasks that can be undertaken by 'outdated' hardware, particularly when it comes to servers. I hope a lot of readers did find that tutorial useful.

No news is bad news

I felt compelled to write and thank you for the more comprehensive news section (issue 9), which was much more informative than in previous issues.

I myself have been using Linux for a number of years, and various other shades of Unix before that, so I find that a lot of your tutorials are teaching to the converted. The reviews I do find interesting, and the occasional feature (e.g. *Mozilla*), and also the cover disc sometimes.

Some people may say you can get all your news online these days – and that is true if you want to spend all day trawling various websites, most of which mix up news and comment and give everything the same priority. I don't really have time to do this, and I find your coverage better than some of the 'professional' magazines, who seem to merely announce new product releases.

On the subject of features, though, could we have more case studies and/or profiles of open source initiatives that might be interesting to those of us who run Linux on a professional level?

J.M. Kinnear, via email

Well, I'm glad you like the new look news. I think that the advantage that magazine news has over websites is that it is easier to prioritise the stories, and filter out some of the irrelevant stuff. If you do want to check out more community based news, we do have regular updates on our own website (www.linuxformat.co.uk)

In terms of case studies, we are looking at including more of these in the mag, as they can often be of benefit to a wide range of users, and there are many interesting Linux stories out there waiting to be told. We do try to make the magazine have something for everyone.

Debian Amiga

Yay! Linux for Amiga! I was really pleased to find that a Linux distro for the Amiga would finally be on the CD. I'm greatly pleased not just that you finally did it, but that you got me a version with a screen mode option. I have had

One reader in California wrote in to tell us how much he enjoyed seeing this picture in the mag - so here it is again...

problems with all the alternative OSs I previously got hold of – Linux used a screen mode that made no sense to my monitor at all, while both the BSDs used NTSC mode. (It could handle the last with a bit of twiddling of the vertical hold, but that didn't really help because all I found out after getting the display legible was that one version had a vital part missing, while on the other version eight of the ten gzip files were suffering from CRC errors) I hope this version is complete and intact.

Changing the subject rather abruptly, I can't help wondering what "horrendous infrostructure" (sic) Roger Fearnside is referring to. He's looking for a file (which he apparently forgot the name of after it had been on the screen for four hours), so I assume he means the directory tree. Sure, things have different names and are not found in the familiar places (mostly because familiar places like C:\Windows don't exist), but at the end of the day it's a directory tree like any other. It just doesn't have an icon called My Computer on top.

BTW, that's one thing you should not try to think too hard about. The first thing you see on a Windoze PC is a 'Desktop' with an icon representing the computer it's on. Clicking on this opens a window with all your

storage devices in, including the main hard disk on which you find a directory called Desktop representing the desktop you started on. Windows isn't an operating system, it's a Klein bottle! ;-)

Ragnar Fyri, via email

The Debian version was just a base install, but it should certainly enable you to get Linux up and running, after which you can take advantage of the Internet or our CDs to add to your collection of applications.

You are right in that some people view the Windows file structure as being the correct one, and any deviance from this is madness. There are pros and cons to everything, and one of the pros of Linux is the ability to logically add filesystems to any part of the tree. Windows arbitrarily assigns drive letters to the devices available as it comes across them, which can cause all sorts of problems. Maybe some people find the structure easier to understand, but it's probably just because it's an effort to learn new things, rather than any inherent 'easiness' of one system or another.

Web community

For the last year I've been working on connecting my SuSE Linux box to the

Internet via a standard dialup account with Demon. 'So what!' you may cry. Well, the point is I've managed to set up a small, but workable web hosting system without any assistance or intervention from Demon.

There's nothing revolutionary about the technology, broadly the system uses: A fixed IP address (a rare commodity now, I know), BT Business Highway/Surtime, Demon's Premier Connect Plus, Easyspace's domain name redirection service, DNS and *Apache* suitably configured for virtual hosts.

The net result of this is effectively a 64k leased line for the cost of a dial up account and BT's Surtime/Business Highway. This still weighs in at around £390 per month though, and so I've decided to share my server with other small business users. I believe that my link will sustain around 10 – 15 small business sites with hit rates up to 500 each per day. These small businesses stand to save a significant amount on the cost of a traditional hosting service and, furthermore benefit from the bespoke service that I can offer to a small number of clients, (Needless to say, hosting isn't my day job).

To promote my idea, I'm calling Cybersharing. I've set up a free, 'dating agency' at www.hippogroup.com to connect owners of web servers, like my own, with website owners seeking a host. I hope you'll take a look at my site and give it a plug, as I have no direct, financial interest in this service – apart from having my own advert listed there!

Andy Chessum

The hippogroup 01284 386848

So basically you just want a free plug in the magazine? Outrageous...



LINUX FORMAT

To have your say

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

What on earth is CUPS?

Do you think printing with Linux is stuck in the dark ages? Then you should join **Nick Veitch** as he gets in his CUPS...

Printing with Linux, and other UNIX systems, has always suffered from a certain lack of functionality. While the old spooling systems were fine for academic network setups where everyone shared a common daisywheel, time – and printers – have moved on. And now Linux printing has moved on to, thanks to CUPS.

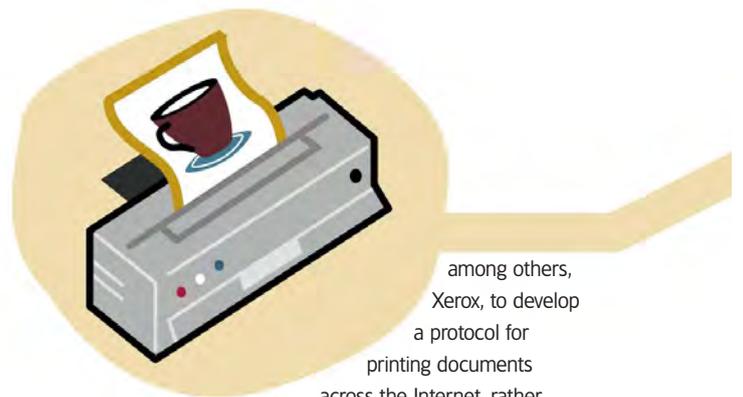
OK, WHAT ON EARTH IS CUPS THEN?

CUPS stands for Common Unix Printing System. It's designed to be a portable printing layer for UNIX-based operating systems, using the IPP protocol to manage jobs and queues...

WOAH THERE! PORTABLE PRINTING LAYER? IPP?

This just means that CUPS is designed to work across a range of UNIX-like systems, including Linux. So even if you are in a mix and match network environment, with, for example, Solaris and Linux boxes, you can use CUPS transparently across both.

IPP stands for Internet Printing Protocol, which is a new standard for printing being supported by the Printers Working Group, though it was really started off by several groups working independently. Back in 1996, Novell co-ordinated a working group, including



among others, Xerox, to develop a protocol for printing documents

across the Internet, rather than just through local networks. This became known as the Lightweight Document Printing Application project. At around the same time, IBM had been working on their own web-based printing technology, HTTP (Hyper Text Printing Protocol). Various other people were also known to be working on similar projects.

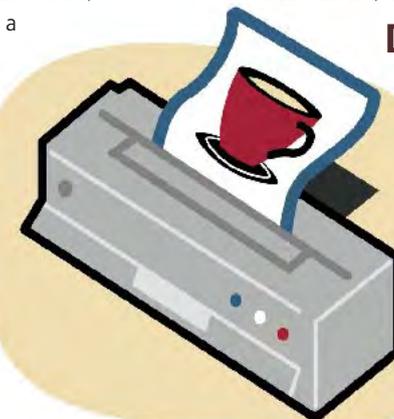
In order to standardise this work, the projects were then co-ordinated under the guidance of the Printer Working Group, a body comprised of people from the printer technology and sales fields.

ISN'T IPP JUST ANOTHER CONFLICTING STANDARD?

Well, it's true that it is another standard, but it is supposed to harmonise the printer industry, with benefits to both the manufacturers (who won't have to build printers that support 20 different protocols), developers (who won't have to build printer interfaces for all the protocols), and consumers (who will, hopefully, be able to buy a printer and expect it to work with the maximum amount of connectivity and support from the rest of their network, hardware and software).

DON'T JETDIRECT AND LPD WORK OVER THE INTERNET?

Yes they do, but neither of them provides proper authentication services. This means that anybody can use them – with IPP and CUPS you can make sure



only recognised users and groups can use your printer. The other protocols don't have any way of checking what facilities the actual output device has either, which can make it difficult to get the kind of output you want. IPP supports pretty much all of the supported printers options, so you know whether it is capable of, for example, 600dpi output before you start.

SO I NEED AN IPP CAPABLE PRINTER?

No, as long as the server which the printer connected to supports IPP, like CUPS. But there is more to CUPS than a system for printing to IPP printers, you can use it for local printers, remote Unix print servers (i.e. under lpd) or even Samba print servers.

WHAT ABOUT PRINTING ACROSS AN APPLE TALK NETWORK?

Yes, you can do that too. As CUPS is open source and has loads of documentation, it's easy to create new backends for the system,

allowing it to print across all sorts of networks, and to all sorts of printers, if people write the drivers.

SO I'LL HAVE TO CODE IT MYSELF THEN?

No, because somebody already has. Helge Blischke has written a backend Perl script for CUPS, which uses *netatalk* to print over an apple network (so you'll need to install *netatalk* too. You can check out our *netatalk* tutorial from issue 1 (if you have it – sorry, there are no back issues left) or check out www.thehamptons/anders/netatalk for a guide on installing it from Anders Brownworth. The PAP Perl script can be found at www.cups.org/bazaar. Though if you do want to code a driver yourself, it isn't that complicated. The *netatalk* one is just a short Perl script.

THIS IS ALL VERY WELL, BUT I HAVE UNIX NETWORK PRINTING ALREADY...

You may well have, but I bet it isn't as easy to use from the administration point of view. CUPS will run on multiple versions of Linux and Unix transparently. Using traditional network printing solutions, this requires a lot of manual configuration on

each client. But CUPS is smart, and servers can broadcast connected printers to all clients on the network, making printer administration much simpler. You connect a printer to one server, and all the clients will see it.

CUPS will run on multiple versions of Linux and Unix transparently

CUPS can also intelligently administer classes or groups of printers, and allow users to print on the first printer that becomes available. Because you can define the same printer on multiple servers, you can easily load balance servers and

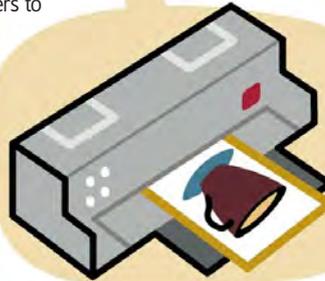
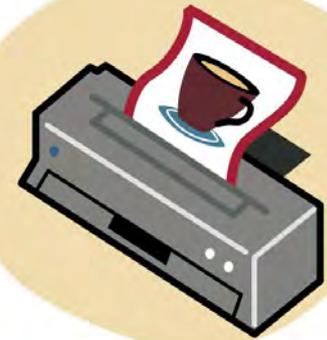
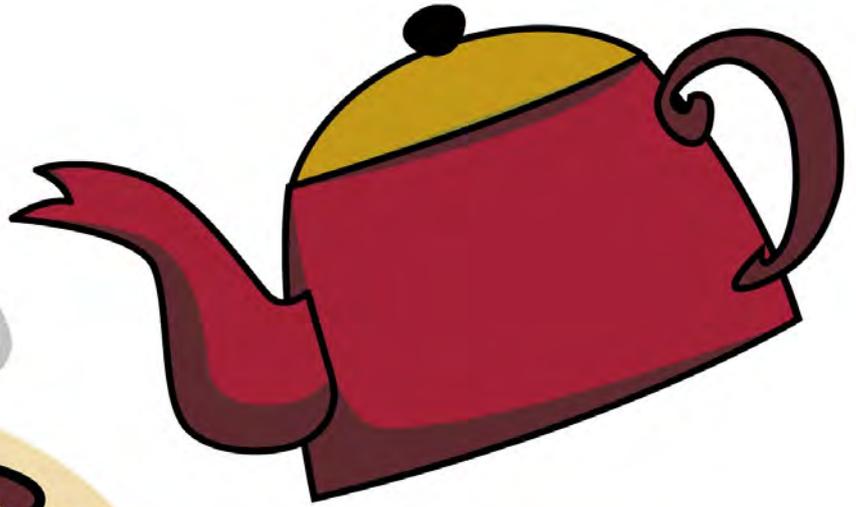
have a fail-safe should one need to go down.

AH, BUT SOME OF MY SYSTEMS ARE INEXTRICABLY LINKED TO LPD DEMONS

No problems, because the latest version of CUPS (1.1) supports lpd clients too, so they should still work without any difficulties.

AND DOES CUPS WORK FOR SAMBA?

Yes. You will need at least version 2.0.6 of Samba running for it to work properly. All you have to do is install CUPS, then edit the *samba.conf* file to include the line **printing = cups** in place of whatever printing commands may already be set up. →



WHAT ON EARTH cups

→ IS CUPS JUST FOR NETWORKS THEN?

No, not at all. It will work just as well on a local system. And you'll still be able to use all the other great features of CUPS, like the browser-based management tools for one, which are quite handy even if you are the only person on the network. These provide a nice user interface for things like checking up on your print queues, altering output specifications and even configuring your printer in the first place

HOW DO I USE THEM?

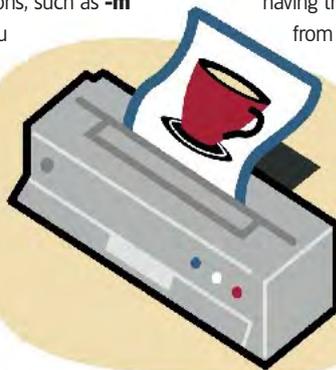
You just need a web browser. Open it up and point it at <http://localhost:631>. This is the port used by IPP by default. Then you'll be able to check up on what your printer is doing. You can use command line tools as well, but the web interface is probably easier. It will even run under *Lynx* fairly well, so don't worry if you aren't running X.

WHAT SOFTWARE SUPPORTS CUPS?

CUPS replaces the `lp` command, so software and scripts which spool their print files this way will still work. You can still use the `lp` command yourself to send files, though it now has a few more options, such as `-m` which will mail you when the job is completed.

SO EVERY PIECE OF SOFTWARE I HAVE WILL BE ABLE TO USE CUPS?

Broadly speaking, you should be able to get everything to work with CUPS, because almost all software supports, at the very



WHAT ABOUT POSTSCRIPT FILES, TEXT DOCUMENTS ETC?

You can print those through CUPS too, fairly easily. You can of course, use the `lp` command again, or you can use one of the many GUI frontends that have appeared, like *xpp*.

Xpp stands for X Printing Panel, and basically is a graphical interface for sending files to the printer. You can set an awful lot of print characteristics here too, such as which tray to use, DPI settings, and some specific features of supported printers such as JCL Economode. It's also possible to adjust things like gamma levels, scaling, image saturation and text mode options.

You can find *xpp* at cups.sourceforge.net/xpp. There are plenty of other graphical front ends, including QTCups, GtkLP and KUPs, which use QT, GTK and KDE respectively, all supporting similar options. You'll find links to these projects at the cups.sourceforge.net site.

BUT HOW CAN I BE SURE THAT ALL THE EXTRA FEATURES WORK ON MY PRINTER?

Because CUPS only displays the ones that do. Aside from having the printer driver, CUPS can also receive information from the printer about options it does or doesn't support, such as extra print trays, memory and so on.

WHAT ABOUT ADVANCED STUFF LIKE ICC PROFILE SUPPORT?

That's one thing that isn't currently handled by CUPS, though it is on their "To Do" list, along with things such as TLS encryption, LDAP support and more.

YOU HAVEN'T SAID MUCH ABOUT PRINT DRIVERS SO FAR.

There are actually plenty of drivers available for CUPS. The program itself includes drivers for HP and Epson printers, although these tend to be fairly generic (e.g. there is one driver for the range of HP Deskjets, which obviously doesn't take into account all the subtle nuances and options in, say, an HP LaserJet 5M).

If you've got a fancy colour inkjet printer, you can also use the rather excellent *GIMP Print* (gimp-print.sourceforge.net). This driver system was originally written for *GIMP* and *Ghostscript*, but the drivers can be used for CUPS too. Most of these drivers focus on colour inkjets such as the Epson Stylus Photo and Stylus Colour range, the Canon BJC range and some HP printers (including more specific Deskjet and Laserjet drivers). Printing with this system does tend to be a little bit slower, but the excellent quality results are well worth it.

MORE ABOUT PRINTERS

To find out more about running printers under Linux, the best place is www.linuxprinting.org. Not only is this the home of the Printing HOWTO, it is also the home of Foomatic and Cups-o-matic. There are many drivers to be found here and all sorts of printing advice and information.

Manufacturers websites usually contain some technical information about their printers. Sadly most commercial sites are

now geared exclusively to direct selling rather than giving you any information, but some of them still have some useful info.

The Adobe website (www.Adobe.com) is a useful site for general information about fonts, the postscript language, colour models and so on.

For reviews of new printers and price comparisons, why not check out www.maximumpc.co.uk

The Linux Printing website (www.linuxprinting.org) is home to the Linux printing HOWTO, and also cups-o-matic, a custom backend for CUPS which will allow you to use one of the many PPD files generated by Foomatic (also housed on this site).

If you have a Postscript printer, you should be able to use the PPD file supplied by the manufacturer.

And if all those options weren't enough, you can go the commercial route and get CUPS bundled with 2300 drivers from Easy Software Products (www.easysw.com/printpro). The drivers cover a wide range of specific printers – there are over 300 Epson drivers alone – in multiple languages).

HOW DO I FIND OUT IF MY PRINTER IS SUPPORTED?

Simply look on the sites above. The Linux printing website has an easy to search database, as does the Easy Software site. If you can't find a driver listed, be sure to search through the Bazaar section on the www.cups.org site. There are some less mainstream printing solutions covered there, as well as info on new CUPS related software.

OKAY, BUT I BET THERE'LL BE ALL SORTS OF PROBLEMS I'LL NEVER FIND THE ANSWERS TO.

Not at all. CUPS is one of the most documented software projects you'll ever have come across. There are comprehensive HTML and PDF files for everything from using the CUPS software to administering a CUPS network. There is even extensive documentation for how CUPS implements IPP, a description of the various associated file formats and comprehensive programmers documentation. All this is in addition to the standard man pages for all the associated command line tools.

BUT WHAT IF I DO HAVE A PROBLEM, AND CAN'T FIND AN ANSWER IN THE DOCUMENTATION?

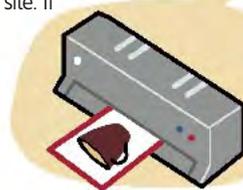
There is limited free support available through the developers of CUPS (www.cups.org/support.html) though of course, they might not be able to help (they do have jobs you know). But if you need commercial support, it is available for the Easy Print Pro software (based on CUPS) – find out more about the package and licenses available at the Easy Software website (www.easysw.com)

OKAY, I'M SOLD – HOW DO I GET CUPS THEN?

The latest versions for a wide range of Unix platforms are available for downloading from the CUPS website. And you will also find the software on our CD in various easy to install formats. If you have Mandrake 7.2 from our Christmas coverdisc, you actually already have CUPS!

If you want to take advantage of the

commercial drivers developed by Easy Software, including professional output devices like linotronic printers, then you should get hold of the Easy Print Pro software from the



www.easysw.com website. Individual licences start at \$95, which includes all 2300 printer drivers.

Unless you have a fairly

standard printer, you'll probably want at least the *GIMP-print* package too. This is also on the CD this month, in source format (no RPMS are currently available for the latest stable version).

AND HOW DO I SET IT UP?

Once you've installed CUPS by whatever method, you can configure it using the command line tools, the web interface or, if you have it installed as part of Mandrake, through the excellent *DrakConf* system. The latter two are probably easiest, as the command line tools will require that you know the name of the driver. You can use the web interface quite successfully with *Lynx*, or another console based browser, so there is no need to even be running X.

Simply go to the web server on your local machine (<http://localhost:631>) and you'll see the main CUPS screen. Click on "Do Administration tasks". You will probably be challenged for a username and password – enter **root** and your root password, then just bang in the details for your printer and you're away.

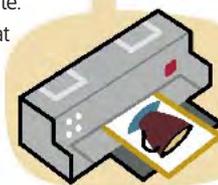
AND IT'S AS SIMPLE AS THAT IS IT?

Yes. It is, trust us. As long as you choose the right drivers for your printers, you shouldn't have any problems at all.

AND HOW WILL I KEEP INFORMED ABOUT NEW DRIVERS AND OTHER DEVELOPMENTS?

Big question, short answer: you can check the website now and then, but you can also join the official CUPS mailing list.

Visit www.cups.org to subscribe. 



CUPS contacts

Michael Sweet
Easy Software Products
44141 Airport View Drive
Suite 204
Hollywood, Maryland
20636-3111 USA
+1 301 373 9600
www.cups.org

Email for cups-related queries: cups-info@cups.org



We've got another fine selection of different tutorials for you this month, from useful one offs to self contained parts of series like our excellent Java tutorial. For those of you with one of those ubiquitous geek-chic accessories, a Palm Pilot, there's a useful guide to co-existence with your desktop Linux box, and we take a detailed look at the NFS filing system, and how to set it up for maximum ease of use! And there's a primer on Perl programming too. Enjoy!

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!

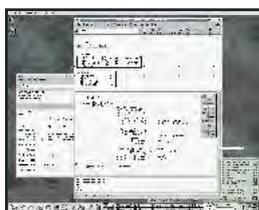
IN THIS ISSUE...

BEGINNING THIS ISSUE

Perl Primer

In the first of our new series about the Perl and its uses in the realworld, Charlie Stross delves into the history of this powerful language, takes a brief look at the developments expected in Perl 6 and demonstrates the anatomy of a basic programme.

p68



PROGRAMMING

Speaking Java

This month, Richard Drummond examines the various ways that Java can exchange data with the outside world...

p73



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

LINUX CONNECTIVITY

Palm and Linux

Palm Pilots (and other PalmOS machines) have been around for a while now and – great news – they're completely at home with your Linux system. This tutorial has everything you need to know to get your diminutive organiser talking to your bulky PC and vice versa. Grab a serial cable and get connected!

p77

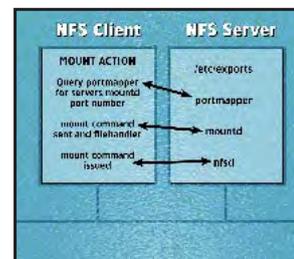


NETWORK FILE SHARING

Network File System

NFS is the default method for sharing files across Unix-style server systems - it's easy to set up, easy to use, and even easy to administer if you know what you're doing! Get the facts here.

p80



DON'T GET OUT OF YOUR DEPTH

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



PERLS

Programmable Extraction and Report Language or Pathologically Eclectic Rubbish Lister? **Charlie Stross** begins a series of tutorials on the realworld applications of Perl.



When people mention programming languages and Linux (or UNIX), the first one that springs to mind is C. C is an efficient, low-level language, and the kernel and most of the standard Linux utilities are written in it. However, C isn't the ideal language for all tasks; in particular, tasks like file manipulation (which the shell is designed for) are cumbersome in C. The shell (or rather, shells) are high level scripting languages designed for gluing command-line programs together; they make it easy to feed a series of data files through filters that slice and dice them, but are far less efficient than C – in particular, just about every time you do something in the shell you end up running an external program (written in C).

BRING IT TOGETHER

The result of this problem (easy to program shells that are inefficient, and C or C++, which are efficient but hard to program) has seen a seemingly endless proliferation of “little languages” that try to bridge the gap between shell scripting and C. For example, a common task is to scan the contents of a file and look for a pattern (using a tool like *grep*), or to modify the contents of a file wherever a search pattern is found (using a tool like *ed* or *sed*). Because it's hard to use *sed* or *grep* to do something more complex than a find/replace operation within a file, the UNIX research team invented *awk*, a pretty little language with C-like syntax, *sed*'s ability to scan through a file and do things whenever it spots a pattern. And thus was born yet another little language.

This trend towards lots of little languages continued until

Perl was born. Perl – the Programmable Extraction and Report Language, or Pathologically Eclectic Rubbish Lister, according to its author, Larry Wall – is ten years old. In the late eighties, Wall had to write tools to support an NSA networking project on VAX minicomputers and Sun workstations. He needed the pattern matching abilities of *awk*, but the networking power of C – but he didn't have time to write it in C. Instead, he invented a new, general purpose language that had C-like syntax, the text processing power of *awk*, string handling routines stolen from BASIC, report

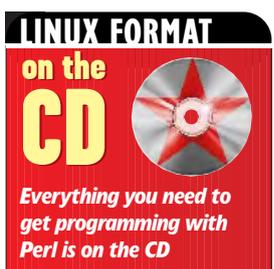
Perl is the language of choice for gluing web applications together

generating routines taken from RPG-2, and a host of other features. Perl wasn't designed to be elegant, it was built to be functional. He released it as freeware in 1988 and it rapidly migrated

to other operating systems; today you can get it on Windows, all flavours of UNIX (including Linux), MacOS, and exotica like IBM mainframes running MVS and Psion Series 5's.

Along the way it has picked up a reputation as the scripting language of choice for gluing web applications together; it integrates particularly well with Apache, the most widely used web server.

Perl is, in fact, one of the first Very High Level Languages (VHLL) – a category that includes Python, Tcl, and Ruby. As the VHLLs have caught on, the rate at which small languages were appearing has dropped off; these days, specific tasks tend to be integrated into an existing VHLL, rather than having a whole new language designed around them.



OF WISDOM

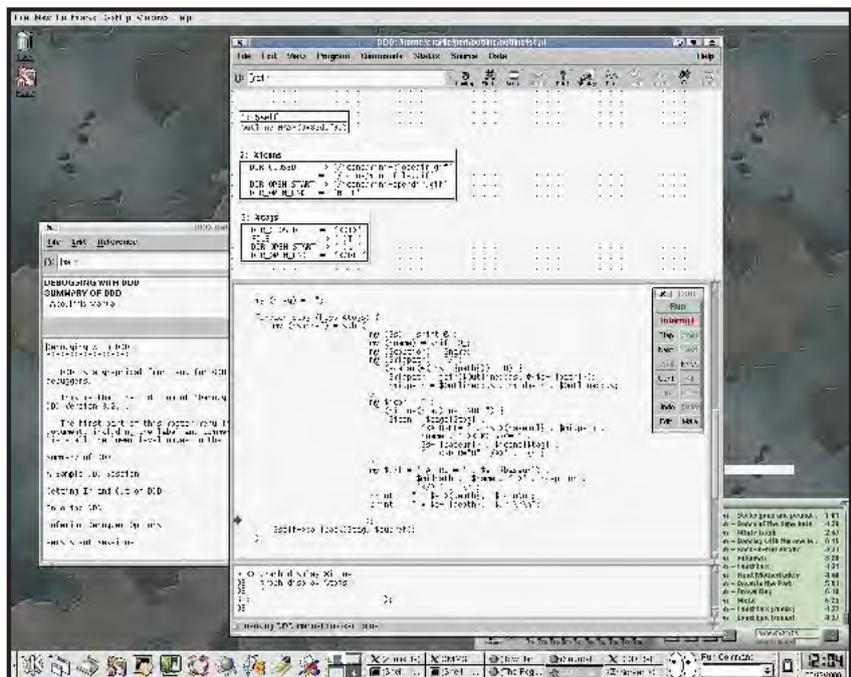
Today, Perl is a vital component of the Linux world. SuSE and Debian Linux won't even install or run without Perl; large chunks of their core system administration code is written in Perl. It's optional on Redhat, but just barely (Redhat picked Python instead); however a fully usable Perl set-up is a standard part of Redhat's default installation. You can locate the Perl interpreter on your system by typing **which perl** – it usually lives in `/usr/bin/perl`.

RUN, PERL, RUN

Perl is currently neither an interpreter nor a compiler; it's a weird hybrid called an interpil. When you type `'perl myprog.pl'`, you are telling Perl to open the file `myprog.pl`, read the contents, and parse them; Perl internally compiles the program and then interprets the resulting data structure at high speed. This way of doing things is considerably faster than a normal interpreter, which parses, tokenises and acts on each line in a program as it reads it, but is a little slower at start-up than a true compiler. On the other hand, it's easier to use at the command prompt. This way of working is likely to change in Perl 6 – of which, more later – that is likely to share Python's way of doing things. Python is a compiler, but you run it like an interpreter. The first time it runs on a file it compiles it to an intermediate file, then executes that; thereafter, it runs the compiled file unless you've updated the source file in the meantime.

Unlike most other programming languages, Perl was designed by a linguist, and Larry Wall has said (frequently) that one of his priorities was to bear in mind that computer languages are written for the convenience of human programmers. So Perl isn't quite as rigid as a more traditional language such as C or Pascal. Of course, the old saw about donating a sufficiency of rope to facilitate a hanging applies: Perl's flexibility can be a problem: the commonest answer to any question about accomplishing some task in Perl is "there's more than one way to do it." In Perl, there's always more than one way to do it!

Perl borrows from other languages. Many of its flow of control constructs resemble C; in other ways, including its standard variable naming scheme (variables are prefixed with a symbol like `$` or `%`, rather than being pre-declared, as in C) resemble those of the UNIX shell. Perl provides a huge



repertoire of built-in keywords, and here it is even more eclectic; just about the whole UNIX system-call library is represented (so that you can resolve an Internet hostname using `gethostbyname()`, or read the current system time and date using `time()`), along with other operators looted from all over the place; BASIC programmers will find familiar `substr()` and `length()` operators for use on strings, for example.

In addition, Perl offers specialised sub-languages for doing common tasks. Regular expressions provide a pattern-matching language for search and replace operations on strings; Perl's regular expression system is massively more powerful than most. Much of Perl's power comes from CPAN, the Combined Perl Archive Network (see <http://www.perl.com/CPAN/>), a huge library of modules for specialised tasks. Among these, probably the most frequently used is the CGI module – which makes writing web CGI scripts easy. But Perl is much more than just a web server scripting tool. It's really a general-purpose application language that is best suited for writing back-office programs: complex data-munging tasks that involve communications, →

The data display debugger hard at work, um, debugging a Perl program.

Running a program

To get a Perl program to run on Linux, you type it into a text editor and save the file. Make sure that the first line of the file begins with the characters `#!/` (omitting the quote marks) followed by the full pathname to the Perl program, for example:

```
#!/usr/bin/perl
```

Then use `chmod(1)` to make the file executable:

```
chmod +x myprog.pl
```

Finally, run it like this, from the current directory:

```
./myprog.pl
```

You can miss out the leading `./` if you have `.` in your `PATH` environment variable.

Alternatively, in the directory you saved the file type: `perl myprog.pl`

→ database access, text processing, and (once in a blue moon) graphical user interfaces.

ANATOMY OF A PERL PROGRAM

Here's a simple Perl program. When you run it, it prompts you to enter some text. It reads whatever you type, until you type in a line containing nothing except the word "quit". It then prints "you typed", followed by a line-numbered listing of whatever you fed it (before the "quit" line).

```
1: #!/usr/bin/perl
2:
3: $in = "";
4: @lines = ();
5:
6: print "Enter some text >";
7: while($in = <STDIN>) {
8:   print "\n";
9:   last if ($in =~ /^quit$/i);
10:  push(@lines, $in);
11:  print "Enter some text >";
12: }
13: chomp @lines;
14: print "\n";
15: print "You typed:\n";
16:
17: for ($i = 0; $i <= scalar(@lines); $i++) {
18:   print sprintf("%03d", $i), $lines[$i], "\n";
19: }
20:
21: exit;
```

HERE'S HOW IT WORKS

The first line is optional but makes life easier; it's not strictly part of the program, but it tells the shell that this is a Perl program and you run it by invoking `/usr/bin/perl` and feeding the rest of the file to that program. (If your Perl interpreter lives somewhere else – say in `/bin/perl` or `/usr/local/bin/perl` – you need to change this line accordingly.)

Next, we initialise two variables: `$in` and `@lines`. These are different types of variable.

Where C or Pascal have low-level data types such as integers, floating point numbers, and character types, Perl has three high-level data types: singular, plural, and dictionary.

Line 6 should be fairly clear; we prompt the user to type something. Line 7, however, is a lot more complex. We're introducing the while loop, a flow-of-control construct that looks rather like its equivalent in C; the bracketed block is executed while the expression at the top evaluates to true. However, the expression looks like nothing in C; what we're doing is using some Perl magic to read a line from the standard input file handle `STDIN` (whatever you type on the console) into a variable called `$in`. If the standard input closes the loop exits; otherwise we keep running continuously... except that on line 9 we have two more bits of Perl magic.

Perl's **while** loop has some seriously powerful exit conditions. **last** is a keyword that means 'exit this loop now'; because Perl is happy with somewhat colloquial usage, you can put the if statement controlling this command after it. Basically, line 9 means 'exit this loop if the variable `$in` matches the pattern `/^quit$/i`'.

The `//` actually acts as brackets around a regular expression. A regular expression is a pattern that is tested against the variable preceding it (which it is bound to by the `==` operator). If the pattern matches the variable on the left, the `==` operator returns 'true'; otherwise it returns 'false'. The pattern matches if the beginning variable (in the pattern, `^`) is followed by the characters "quit" and then the end of the variable (in the pattern, `$`), in a case-insensitive manner (that trailing letter `i`.) So if `$in` contains nothing but 'quit' (because the pattern is anchored at either end to the beginning and end of whatever it's matching), the expression returns 'true'... and causes the **last** command to be executed, breaking out of the loop.

Line 10 introduces a new command: **push**. Push takes at least two parameters – an array and one or more scalars. It pushes the scalars onto the front of the array, shuffling everything already in the array to the right to make room. We use it to stash the lines we read in the array `@lines`.

Line 13 shows the **chomp** command in action. Chomp takes one or more scalars (or an array) and strips off any trailing carriage returns off the end of them. We can add them back later; when you put the escape symbol `\n` in a double-quoted string, it is replaced by a carriage return – just like in a C `printf` statement.

Finally, we get to lines 17 to 19. These show us the format of a Perl **for loop**; it's just like a for loop in C. Here, we use it to increment a counter called `$i` (a scalar containing an integer number) and, using **sprintf** (which works pretty much like it's C equivalent) to print each line of the array `@lines` in turn, prefixed with the line number. The one point to note is that to refer to the fifth scalar element of an array called

`@fred`, you talk about **\$fred[5]**; the leading dollar sign means 'interpret this in a scalar context'. If we talked about `@fred[5]`, we would actually be using list

context: it would mean 'the array consisting of the fifth element of `@fred`', not 'the scalar which is the fifth element of `@fred`'. This is an important distinction as we'll see in future articles: as I said earlier, Perl variables are singular (scalar), plural (array or list), and dictionary (hashes, aka associative arrays), and most Perl commands do different things in an array or scalar context.

WHAT'S PERL 6 ABOUT?

Perl 5 has been around since 1995, and for some time now it's been showing signs of age. Following a series of upgrades (it's now at release 5.6) the Perl 5 source code is difficult to extend. For that reason, in the Summer of 2000 Larry Wall announced a project to develop a new version of Perl. →

Where C and Pascal have low-level data types, Perl has three high-level types

VARIABLES IN PERL

A scalar (singular, unitary) variable has a name beginning with a dollar symbol, like \$in, and stores one entity – a number or string or just about anything else. Scalars are elastic – they stretch to fit whatever you put in them, so that you can cram a very large binary file into one if you want.

The most important characteristic of a scalar is that it is atomic – that is, Perl can't subdivide it. You can dink with the contents of a scalar using Perl's string or arithmetic operators, but the scalar itself is a single identifier that refers to a single object.

In contrast, when you've got more than one item you can store the data together in an array: these have names beginning with an @ symbol, such as @lines.

Arrays are dynamic; you don't need to specify how many elements there are in one, they grow or shrink to fit. An array consists of a numbered list of scalars, starting from zero: that is, the first item in the array is a scalar with index number (subscript) of zero, the second is a scalar with subscript one, and so on. One oddity to note: when you're referring to a scalar member of an array, you prefix it with a dollar sign to indicate it's a scalar:

```
@my_array = ($foo, "blue", 1, $bar);
print $my_array[1], "\n";
==> blue
```

We have created an array here by assigning a list to it. A list is just a collection of scalars, delimited with brackets; it won't be treated as a separate variable until we explicitly turn it into an array.

We can extract slices from an array like this:

```
@small_slice = @my_array[1..2];
print $small_slice[0], " ", $small_slice[1],
"\n";
==> blue 1
```

@small_slice is a new array; we assign to it a slice of @my_array containing the elements with subscripts between 1 and 2 (inclusive).

In the first print example, we observed that you prefix a scalar member of an array with a dollar sign, to indicate it's a scalar. When printing a slice of an array, you use an @ sign instead of a dollar:

```
print @my_array[1]
```

means "print the array slice containing elements with subscripts from 1 to 1" rather than "print the scalar with array subscript 1". Because the print command understands arrays, this will work and produce exactly the same output as printing the scalar – but not all Perl commands are this forgiving. It's a good idea to get used to distinguishing array slices from scalars as soon as possible, otherwise you're liable to trip over them later.

It's important to note that lists and arrays are not the same thing in Perl. An array is a special type of variable that contains one or more scalars, but can be manipulated as a variable. A list is just a grab-bag of scalars (or other variables). This has far-reaching and subtle consequences. For example, we can force an array to grow by referring to an element in it with a high subscript

```
@big_array = ("foo", "bar", 53, $blue);
# @big_array now contains four elements
$big_array[72] = "";
# We just assigned the empty string to the
# 73rd element of @big_array,
# implicitly populating elements 4..71 with
# empty scalars
```

Just as we have singular (scalar) and plural (array/list) variables, most Perl commands operate differently in a singular (scalar) or plural (list) context. Context is unique to Perl: most languages, such as C or Pascal, insist that a function can only operate on one type of input data. Perl does things differently. For example:

```
print @my_array;
print $my_array[0], $my_array[1],
$my_array[2], $my_array[3];
```

These uses of the "print" function both do what you'd expect even though they're working on different data types, because the print command normally works on a list of one or more scalars; it interprets its parameters as a list (in "list context", in Perl jargon).

If you want to know how many cells @my_array has, use the function scalar() to force it into a scalar context; when you refer to a list or array in scalar context, Perl returns its size.

```
print scalar(@my_array);
==> 4
```

There is one other type of simple variable in Perl: the hash (or associative array, or dictionary). This type of variable is found in other high-level languages but not in the likes of C, Pascal or Basic. A hash is basically an array, but instead of each scalar element of the array having a numerical subscript starting with 0, 1, 2..., the hash subscripts are strings (or numbers, or anything else that's legal as the value of a scalar). Hashes are indicated by a % sign

```
%dictionary = ( "my_colour" => "blue",
               "my_side" =>
$direction,
               "num_of_widgets" => 4
               );
print $dictionary{my_colour}, "\n";
=> blue
```

Each entry in a hash consists of a key and a value; you use the key where you'd use a subscript in an array, and the value is some scalar entity. In fact, we can assign an array to a hash, as long as the array has an even number of contents so that it can be broken down into matching key/value pairs. In the example above, we create %dictionary by assigning a list of key/value pairs to it; the => operator is a synonym for the comma (,) in Perl which makes it easier to see that the key (my_colour) points to the value ("blue").

There are two special commands to remember with hashes: keys() and each(). keys(%dictionary) will return a list of all the keys in %dictionary, i.e. a list like ("my_colour", "my_side", "num_of_widgets"). This makes it easy for us to loop all the items in a hash and do something with them. In scalar context, keys() returns the number of elements in the hash.

If the hash is huge, the list of keys will be huge, too: so Perl also provides the each() command. Whenever each(%dictionary) is called, it returns a two-element list consisting of the next key, and the value associated with it – until the end of the hash is reached, at which point it returns an empty list. So we can also walk through the elements of a hash using a while() loop

```
while (($key, $value) = each (%dictionary)) {
# do something with $key and $value -
leave the loop if these items
# are empty because we've reached the
end of %dictionary }
```

→ The Perl 6 process is a mixture of public participation and benign dictatorship. First, there was an RFC (request for comment) process – users contributed several hundred feature requests for language extensions. Wall is now sorting and ranking these, before making a final determination of what will go into Perl 6. In the next stage, a group of developers will take over various sections of the new code.

One goal of Perl 6 is that it must run 95% of Perl 5 programs unchanged, and that of the remaining 5%, the changes required should be minimal. So almost everything in this tutorial should work for Perl 6 as well. (However, we're told that typeglobs are definitely going to go away.)

Perl 6 is going to overhaul both the language itself, and the back end compiler and interpreter. It's likely to work more

like Python, where scripts are compiled to a bytecode format which can then be executed rapidly. The compiler will be able to emit C, Java, and C# code as well, and if you invoke it on an already-compiled script the bytecode won't be recompiled (thus reducing startup time).

One goal is for the language to become dynamically extensible, using Perl modules to implement mini-languages; non-critical commands (such as those for socket handling) will be exported into dynamically loaded code. There's going to be a complete overhaul of Perl's built-in data types and object-oriented features, and Perl will acquire function and method signatures (replacing the current broken prototype and parameter passing mechanisms) to allow better compile-time checking.

What all this means is that Perl 6 is intended to be a meta-language for creating domain-specific languages. With support for functional programming and logic programming as well as OOP, it's going to become ever more flexible: and with mature compiler support and function signatures it's going to be more useful in large-scale software development projects.

The first betas of Perl 6 should emerge – with luck – in late Summer, 2001... **LXF**

NEXT MONTH

Next month, **Charlie Stross** continues our tutorial with a look at namespaces, subroutines, and how to write large programs in Perl. And if that's a bit basic for you, he talks about `eval()` and closures – how to make Perl write Perl on the fly!

PERL ON THE WEB AND ON PAPER

Perl culture has its roots in two places; the Internet, and publishing house O'Reilly and Associates Inc (<http://www.oreilly.com/>). O'Reilly and Associates fund Perl's inventor, Larry Wall, to work full-time on his brainchild; they also publish the best and broadest range of books about the language. In fact, no Perl programmer should be without a copy of the definitive book on the language, *Programming Perl* (third edition), by Larry Wall, Tom Christiansen, and Jon Orwant (ISBN 1-56592-00027-8). If you're taking your first steps with the language, you should also consider *Learning Perl* by Randal Schwartz (ISBN 1-56592-042-2) – this was the first tutorial on the language, and there's no way that a magazine column such as this one can compete with it. You may also want to buy *The Perl CD Bookshelf* (ISBN 1-56592-462-2), a CD-ROM which comes with a paper copy of *Perl in a Nutshell* (a quick-reference guide that's only 600 pages long!) and both of the above books, along with three others.

Finally, if you know Perl but can't remember every little detail of the language syntax, you need a copy of the *Perl 5 Pocket Reference* (3rd edition) by Johan Vromans – it's a very terse aid-memoire that fits in just 90 pages, and it'll set you back a fiver or so (ISBN: 0-596-00032-4).

In addition to printing books, O'Reilly maintain (and fund) the core website, <http://www.perl.com/>. This is effectively a portal site for Perl programmers, and has all the documentation online, along with articles about Perl. It also contains a central link into CPAN, the Combined

Perl Archive Network, source of almost all Perl downloadables. (CPAN contains stacks of Perl scripts, the language itself in numerous versions, and a gigantic archive of modules – reusable object classes.)

In contrast, <http://www.perl.org/> is the public

Programming Perl is the definitive text book on the subject.



www.perl.com should be your first stop on the web for all things Perl related.

headquarters of Perl Mongers – a grass-roots organisation of Perl programmers that acts as an umbrella organisation for local user groups. Check it out – there's probably one near you, and often the easiest way to learn is to pick the brains of someone who's already been through the process. There's also an associated web portal, <http://use.perl.org/> which contains news of active Perl events and module releases, and another portal for questions, answers, and support: <http://perlmonks.org/>.

If you're really interested you might want to subscribe to the *Perl Journal*. This is a quarterly magazine for Perl programmers; subscribers can also access its content via the web at <http://www.tpj.com/>. TPJ is authoritative and chock-full of technical articles, albeit

targeting the jobbing Perl programmer rather than the general enthusiast.

Traditionally, most Perl support is delivered via usenet. A number of newsgroups (conferences) exist specifically for Perl, starting with `news:comp.lang.perl.announce` for announcements, then `news:comp.lang.perl.modules` (discussion of object-oriented programming in Perl), `news:comp.lang.perl.cgi` (discussion of Web application programming in Perl, not for general Perl questions!), and `news:comp.lang.perl.moderated` (for questions to which you can't find an answer in the FAQs). Unfortunately `news:comp.lang.perl.misc` is so badly overrun that it's difficult to use it these days.



SPEAKING JAVA

Part 5: Exceptions and Streams

Java is blessed with a rich set of input and output services. **Richard Drummond** takes a look at how they work.

Computer programs are useless unless they can exchange data with the outside world, so any programming language must provide a mechanism for input and output. Not surprisingly, Java is equipped with a rich set of I/O services via its standard class library, and in this, the latest installment of our whistle-stop tour of the language, we are going to examine some of these. Before we can do that, however, we need to take a look at exceptions.

TO ERR IS HUMAN

To build robust software, you must anticipate the sort of errors that will occur during execution and design your programs to check and deal with them. Untrapped errors can cause a loss of data and frustration. In languages like C, there is no built-in support for handling errors – it is up to the programmer to implement a mechanism for flagging and dealing with them. This is not so with Java and its exception handling.

An exception in Java is a method of signalling that some special or error condition has occurred. By generating or throwing an exception, a program can communicate – outside of the normal flow of the program – that something has gone wrong. There is no need to pass error codes as the return values of methods or for their clients to explicitly check for these – as you would in C. The exception mechanism allows you to completely separate the error checking and handling code from the normal, algorithmic path of execution.

When an exception is thrown, the program will jump out of its current code block and look for an exception handler – a special section of code designed to catch exceptions – at that block's level. If no exception handler is found at the same level, it tries at the next level up, and so on, all the way up to the method that the code is part of. If no handler is found

within the method, one is searched for at the point the method was invoked in whatever routine called it and so on – bubbling up through the call stack, until an exception handler is found. If the exception is not caught anywhere, the program will exit with an error message.

Exceptions are generated in Java with the **throw** keyword and caught with **catch**. You specify that you want to catch exceptions in some block of code by wrapping it up in a **try-catch** statement. Have a look at this example:

```
class ExceptionDemo1
{
    public static void main( String args[] )
    {
        int x=5; int y=0;

        System.out.println( "Dividing "+x+" by "+y+"." );
        System.out.println( x/y );

        System.out.println( "This won't be executed." );
    }
}
```

Here we are trying to divide a number by zero. We use variables, rather than using a constant expression such as 5/0, because the Java compiler would try to evaluate such a constant and fail with a division by zero error. When this example is compiled and executed, an exception will be thrown, the program terminated and a rather unpleasant error message dumped to the console. Hardly user-friendly! →



→ Now have a look at this example.

```
class ExceptionDemo2
{
    public static void main( String args[] )
    {
        int x=5, y=0;

        System.out.println( "Dividing "+x+" by "+y+"." );
        try
        {
            System.out.println( x/y );
        }
        catch( ArithmeticException e )
        {
            System.out.println( "Sorry, I couldn't do that." );
        }
        System.out.println( "Carrying on . . ." );
    }
}
```

Here we enclose the offending statement in a **try** clause. Now when the exception is thrown, execution will jump to the following **catch** clause where we trap the exception. There's not much we can do to recover from this error, so we just print a friendly message. When that's done, execution will carry on after the **try-catch** statement.

TO CATCH DIVINE

An exception in Java is actually an instance of the **java.lang** package's **Exception** class – or one of its sub-classes: the **java.lang** package and the others specialise the base **Exception** class into particular types of exception – such as **ArrayIndexOutOfBoundsException**, **IOException**, etc.

To generate an exception, a new **Exception** instance must be created before it can be thrown with the **throw** keyword. This object is passed as an argument to the closest exception handler. The **catch** clause acts like a method and takes an object of **Exception** class as a parameter; like regular methods, it too may be overloaded. That is, you can specify multiple catch clauses to trap different classes of exception. The sequence of clauses will be searched in order until first one that matches is found.

In the above example, we were only looking for exceptions of **ArithmeticException** class (or any of its subclasses if it had any). If some other type of exception were to occur, it would not be caught and the program would terminate as before. To rectify this, we could add

```
catch( Exception e )
{
    System.out.println( "Error:"+e );
    return;
}
```

to the example above after the existing **catch** clause. Then any other exception that might be thrown would be caught. This time we simply print out its corresponding error message

(the **Exception** class's **toString** method is overridden to return its error message as a **String**) and exit the class with a **return** statement. Note, that a sequence of **catch** clauses must go from specific to general; otherwise the general case would trap all the exceptions and the clauses trapping the more specific exception types would never be used.

CHECKED VS UNCHECKED

As well as being divided into class, exceptions in Java are also notionally categorised as checked or unchecked. All unchecked exceptions are instances of the **RuntimeException** class or any of its children; any other exceptions are checked exceptions. Runtime exceptions are any of a variety of exceptions generated by the Java virtual machine itself, and include things like arithmetic errors, out of bounds arrays accesses, and security violations. Unchecked exceptions are application errors such as missing files or network errors.

Why the distinction?, you might ask. Well, unchecked exceptions don't have to be caught; checked exceptions must be. In fact, if we invoke a method which could generate a checked exception and we don't provide an error-handler, then the Java compiler will complain. The rationale for this is that checked exceptions are the everyday sort of errors which could occur at any point in program and are typically difficult to recover from, so it would be unreasonable to require that they are always trapped. On the other hand, checked exceptions are the type of thing that must be caught in order to create a well-behaved program.

How do we know which exceptions a particular method might throw? Well, a method will declare any checked exception that it will throw in its signature with the **throws** keyword. For example, the fictional method

```
void myMethod( void ) throws
FileNotFoundException, SecurityException
{
    /* some code here . . . */
}
```

states that it may throw exceptions of class **FileNotFoundException** and **SecurityException**. Any program that invokes this method must be prepared to deal with these exceptions.

Of course, the methods specified in the various classes and packages of the standard Java class library, may also throw checked exceptions. To find out which methods throw what, take a look at the Java API specification. The **java.io** package which will look at next defines a slew of checked exceptions which its methods may throw. Therefore, to use many of Java's IO services, you need to understand exceptions – which is why we looked at these first.

GENTLY DOWN THE STREAM

The basis for all IO in Java – whether it is with the console, with files, pipes or a network socket – is the stream. The **InputStream** class defines an abstract channel where data can flow into your application and the **OutputStream** class, a

Errors

Exceptions are usually a signal that some non-fatal error has occurred. In more drastic circumstances, Errors are thrown. The Error class is the other sibling of the Throwable class, but Errors cannot be caught. You wouldn't be able to do much, if you could catch them. An Error is a sign that something critical has gone wrong, such as the JVM running out of memory. There's not much even a well-designed program can do in these circumstances.

channel where data flows out. These two classes are not very useful by themselves and are specialised with various sub-classes which provide I/O with different kinds of data to different media, but they do provide a core interface for I/O.

The various stream classes provide services to read and write bytes of data to streams. At the base level, **InputStream** supplies an overloaded **read** method to read a single byte (returned as an **int**) or an array of bytes from a stream, **OutputStream** supplies similar methods for writing. Have a look at this example:

```
import java.io.*;

class HelloUser
{
    public static void main( String args[] )
    {
        byte[] name = new byte[20];

        System.out.println( "What is your name?" );

        try
        {
            System.in.read( name );
            System.out.print( "Hello, " );
            System.out.write( name );
        }
        catch( Exception e )
        {
            System.out.println( "Error:" + e );
        }
    }
}
```

First note that we have to import the classes from the **java.io** package that we need; here we just import them all. Anyway, this example prompts the user to enter their name and then tries to read the next twenty bytes from the standard input stream. What will actually happen is that it will allow the user to enter a line of text on the console followed by return and then put the first twenty bytes of it into the array **name** – although this behaviour is somewhat dependent on which JVM you are using.

The program responds with 'Hello' and whatever name the user entered. Note here, that we use the **write** method to send the array of bytes to the output stream rather than trying to use **print**. That's because **print** works with arrays of **chars** rather than arrays of **bytes** and expects these chars to be Unicode encoded; in locales outside of the Western world, these two will not always be equivalent. Note also that we don't explicitly print a carriage return here. That is because if the user entered less than twenty characters then the carriage return that he or she entered will be contained within the array **name** and so be output when the array is written.

For most cases, though, operating on plain streams is just too much hard work. That is why Java has provided **Reader** and **Writer** classes since JDK 1.1. These provide character-based access to streams.

Take a look at this:

```
import java.io.*;

class HelloUser2
{
    public static void main( String args[] )
    {
        String name;
        BufferedReader in = new BufferedReader( new
        InputStreamReader( System.in ) );

        System.out.println( "What is your name?" );

        try
        {
            name = in.readLine();
            System.out.println( "Hello, "+name+"." );
        }
        catch( Exception e )
        {
            System.out.println( "Error:" + e );
        }
    }
}
```

This example performs more or less than same function as the previous one, but does the job far more elegantly. Here we wrap the standard input stream with a **Reader** class and then wrap that with a **BufferedReader** class. The first step allows to read **chars** from the input stream using the default character encoding, and the second step provides a more efficient buffered access to the stream. The latter also furnishes us with **readLine()** method, so that we can read a line of text up to a carriage return from the input and return it as a **String**. This can then be printed with **println()**.

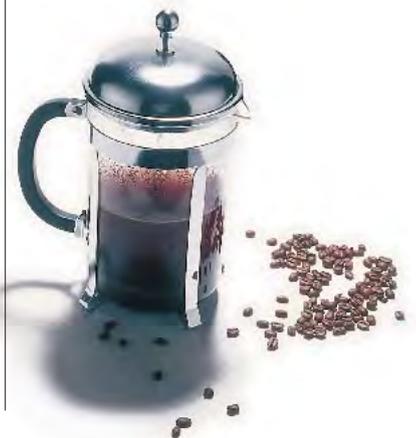
WIDENING THE STREAM

The stream classes in **java.io** don't just allow you to access the standard input and output channels, however. They wouldn't be that useful if they did. The basic stream classes are overridden within the I/O package to handle I/O to files and pipes. Furthermore, the package **java.util.zip** provides stream-based access to files compressed with *gzip*, while **java.net** allows stream-based access to network sockets.

For now, though, we'll just look at files. To access a file for input, you create a new **FileInputStream** object with the pathname of the file to read as parameter to the constructor. To access a file for output, you create a new **FileOutputStream** object. This class has several constructors: one options is to pass the path name of the file to open as a parameter, another is to pass a **File** object, and so on. Both the file stream classes behave like the abstract input and output streams; you can read and write bytes to them or wrap them up as data streams or with reader and writer classes. Likewise **FileReader** and **FileWriter** classes correspond to the reader and writer classes for **InputStream** and provide character-based access to files. →

Streamweaver

Every application has access to three pre-existing streams – the default input, output and errors streams. If the program was started from a console, then typically these refer to the console itself, but they may be redirected by the user to files, pipes and so on in the usual manner. These streams are accessible as public fields of the System class of the java.lang package, namely in, out and err. We have been using the out field all through this tutorial. System.out is actually an object of PrintStream class, an indirect subclass of OutputStream. This class provides formatted printing of strings via the print and println methods and so is indispensable for writing any sort of messages to the console.



Looking ahead

Hopefully from this brief tour I have given you a glimpse of the power and flexibility that Java offers for handling input and output. To find out more, have a look at the Java API documentation at <http://java.sun.com/docs>. The sections on `java.lang` and `java.io` should be fairly digestible to you if you have followed this series.

We have now reached a point in this tutorial series where it should be possible for you to start writing useful, if not wildly exciting, applications. In the next few issues, when we start looking at threads and the AWT (Abstract Windowing Toolkit), the fun will really begin.

→ As a quick example, take a look at this program, an emulation of the shell `cat` command but written in Java. It takes a list of filenames as arguments and dumps each one in turn to its standard output.

```
import java.io.*;

class cat
{
    public static void main( String args[] )
    {
        for( int i=0; i<args.length; i++ )
        {
            try
            {
                BufferedReader in = new BufferedReader( new
                FileReader( args[i] ) );

                try
                {
                    String line;

                    while( ( line=in.readLine() ) != null )
                    {
                        System.out.println( line );
                    }
                }
                catch( Exception e )
                {
                    System.out.println( "Error: "+e );
                }
            }
            catch( Exception e )
            {
                System.out.println( "Couldn't open file
                "+args[i]+"");
                System.out.println( "Error: "+e );
            }
        }
    }
}
```

NOT JUST BYTES AND CHARS

Java doesn't just let you input and output data as text, however. The `DataInputStream` and `DataOutputStream` classes allow the writing of primitive data types such as **booleans**, **ints** and **floats** as raw data. Have a look at this example which writes an array of integers to a file called 'demo.tmp'.

```
import java.io.*;

class IntegerDemo1
{
    public static void main( String[] args )
    {
        int[] primes= { 1,2,3,5,7,11,13,17,19,23 };
    }
}
```

```
try
{
    DataOutputStream out = new DataOutputStream(
    new FileOutputStream( "demo.tmp" ) );

    for( int i=0; i<primes.length; i++ )
    {
        out.writeInt( primes[i] );
    }
}
catch( Exception e )
{
    System.out.println( "Error:"+e );
}
}
```

If you compile and run this example, you can then examine the file generated with the command `hexdump -C demo.tmp`. If you do this you'll see that each integer is stored in the file as four bytes, with the most significant byte first. We can create a simple program to read this file. The next example will read the file 'demo.tmp' and interpret it.

```
import java.io.*;

class IntegerDemo2
{
    public static void main( String[] args )
    {
        try
        {
            DataInputStream in = new DataInputStream(
            new FileInputStream( "demo.tmp" ) );

            while( in.available() > 0 )
            {
                int num = in.readInt();
                System.out.println( num );
            }
        }
        catch( Exception e )
        {
            System.out.println( "Error:"+e );
        }
    }
}
```

What's more, you are not limited to inputting and outputting primitive types with Java. You can even read and write the states of objects to streams with any of the various `ObjectInputStream` and `ObjectOutputStream` classes. Any class which implements the `Serializable` interface can be written to a stream and read back at a later date. This forms a mechanism for persistence – objects can retain their state between executions of a program. In other languages creating persistent objects is very difficult, but Java makes it laughably easy. Yet another reason to use Java. **LXF**



Using your PALM HANDHELD with Linux

Mike Saunders reveals how to keep your PalmOS handheld in order with Linux and a serial cable.

Linux may be the best thing since sliced bread for bulky desktop machines and mighty servers, but if you're looking for computing on the move, then you need to weigh up the other alternatives. Sure, you can buy a laptop and lug it around on your travels, but they're not particularly cheap or portable – and just look at the average battery life on these machines.

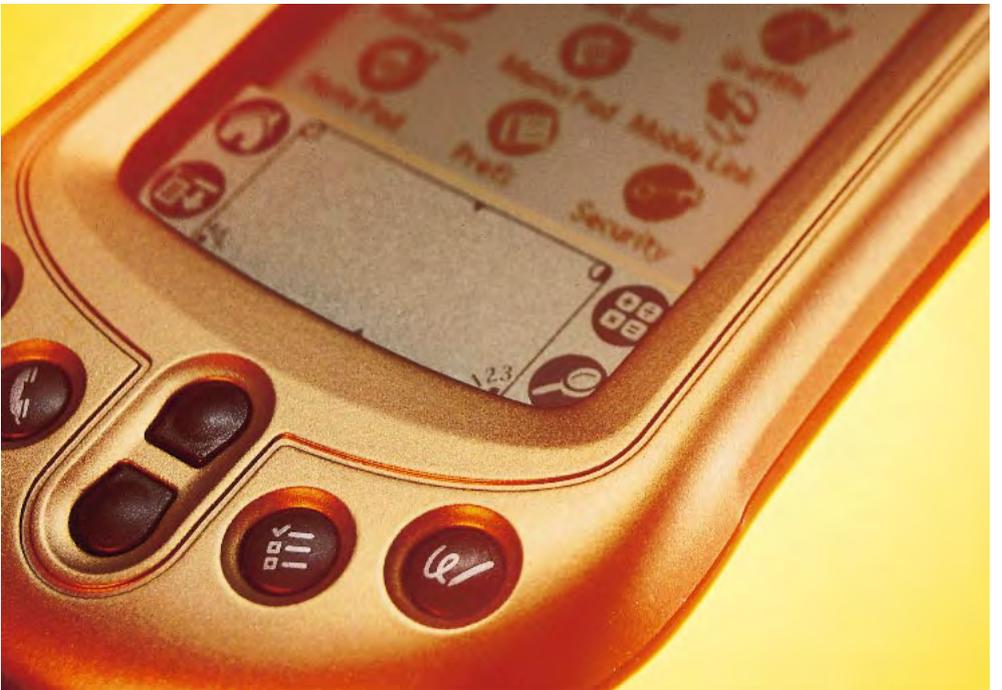
And if it gets stolen, you're up a certain creek without any hint of a paddle. So, in recent years the emphasis has turned to miniature handheld devices which contain the basic functionality you're likely to require, coupled with a diminutive size and increased longevity in their power cells. While not likely to make you trash the home PC, their ever-improving capabilities have proved immensely handy for many.

Microsoft tried their hardest with the Windows CE and Pocket PC platforms, but most of the attention in the past few years has focused on the PalmOS devices. Originally created by 3COM, the Palm Pilot was a simple monochrome handheld organiser with some neat features and tiny dimensions. Since then, many more handhelds running PalmOS have appeared, and the hardware itself has gained memory, a backlit screen, and even colour.

If you're the proud owner of a PalmOS handheld computer with serial cable connection and have Linux up and running on your PC, this tutorial will walk you through the setup process to get your Palm device taking to it (those with USB connections like the Handspring Visor should check out <http://www.milosch.net/visor/first>). We'll also look at some of the applications available to make working with your Palm easier. Are you sitting comfortably?

1. WHAT YOU'LL NEED

Firstly, you may be familiar with the *Palm Desktop* software supplied with your device for Windows systems. This program (both fairly powerful and user-friendly) sets up your handheld



automatically, but don't be daunted by the procedure for Linux. It only takes a few commands before you can be happily transferring data back-and-forth.

The most essential software package you'll need is the pilot-link utilities. These are included on most distribution CDs, so have a search around on the disc you used to install Linux. If it's not there (or you installed over a network), you'll find it at: http://sourceforge.net/project/?group_id=2158. →

RUNNING LINUX ON A PALM

Yes, we know – you probably think this is a joke – but believe it or not, the operating system noted for its performance on large networks and meaty Alpha systems will also run on a tiny little monochrome handheld. Depending on your point of view, it's either the greatest testament to Linux's scalability, or an exercise in porting-gone-mad.

That's not to say that Linux will run straight-off like your standard Mandrake or SuSE distribution. A lot of fiddling and effort is required first, and – unless you're a developer – the end result is little more

than a novelty. However, it's certainly a good sign that Linux has a strong future in the PDA and embedded market, and we're keen to follow the progress here.

Naturally, you can't expect to do a great deal with Linux on Palm devices just yet, but it's interesting to see and – if you have the time and patience – it's worth a quick look.

To find out more about the project which led Linus Torvalds to remark "Whoever came up with the idea must have been on drugs", point your browser at http://mobileminder.com/uclinux_faq/.

→ Users of Mandrake, Definite, SuSE or Red Hat will find the RPM packages (found by searching <http://www.rpmfind.net>) easiest to use. Simply download the “pilot-link.xxx.rpm” file (where xxx is the version), switch to the “root” user at a terminal prompt (using **su**), and enter **rpm -Uvh pilot-link.xxx.rpm**.

If you encounter all sorts of dependency errors, or your distribution simply doesn't use RPM packages, you can build the pilot-link tools from the source code. Download the “pilot-link.xxx.tar.gz” archive, make sure you have the necessary development software on your system (GCC, Make, autoconf etc.), and then extract it with **tar xfvz pilot-link.0.9.3.tar.gz** (or whatever the archive is called). From there, read the README and INSTALL files for more information on compiling and installing the programs. Once this has been completed, you can move onto the next step.

2. SETTING UP

Now that you've installed the *pilot-link* software, you have the essential suite of tools on your system for transferring data between your Palm machine and Linux PC. Of course, life's never that simple and you'll have to do a bit of configuring and setting-up first, but this is very straightforward and quick.

Firstly, you may be familiar with the device system that Linux (and other UNIX variants) use to deal with every part of the system. Each piece of hardware in your machine has a “node” in the /dev directory, which allows programs to talk to these devices.

Now, we need to create a new node that points to the location of the serial port – the place where the link cable between your Palm handheld and the PC plugs in. In fact, there's already a link to the serial port in place, but we'll make a link to it with a more obvious name. Enter the following commands at a prompt (as root) **cd /dev/ ln -s ttyS0 pilot**.

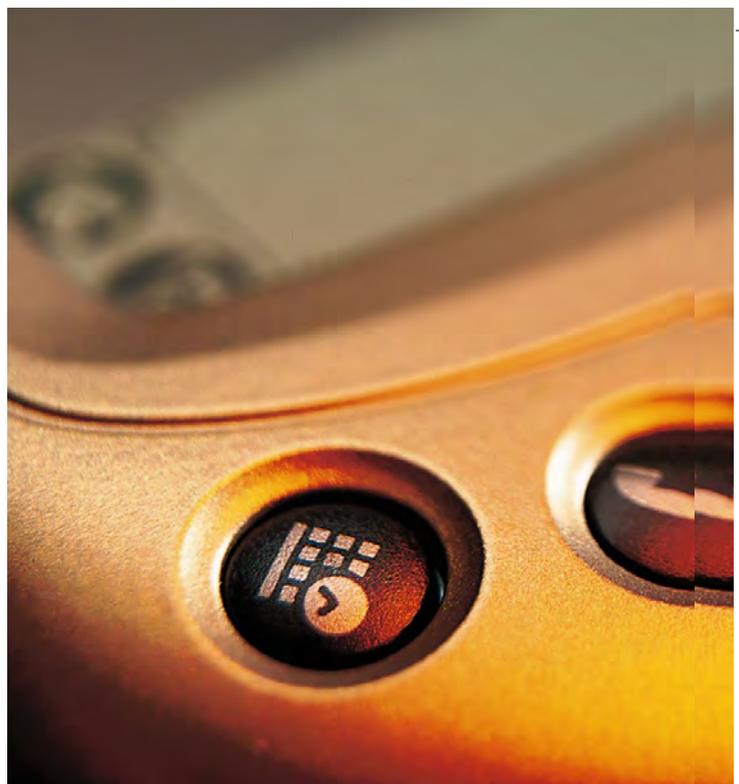
Note here, however, that you will need to change the

number in “ttyS0” to “ttyS1” (or whatever) depending on the port your Palm cable is plugged into. Simply enough, all we have done here is create a link to the “ttyS0” node (which points to the serial port), and called it “pilot”. You may already have a similar link in there called “modem” which links to the same place.

With this sorted, /dev/pilot is now the place programs can look if they want to communicate with your handheld machine. Now we can move onto the process of transferring files.

3 – USING PILOT-LINK

As mentioned before, *pilot-link* provides a suite of tools which perform the most common operations you're likely to need. They're all straightforward



command-line programs which do the job – we'll cover the more comfortable graphical programs later, but it's worth experimenting with these tools.

Quite a few utilities are included in the *pilot-link* package, with programs for reading the address book, installing memos from a text file, adding new software and more (enter **man pilot-link** for detailed documentation). As an example, we'll install a new program onto the Palm.

Download a software package from one of the shareware PalmOS sites, and store it in your home directory. Then, plug the Palm cable into the serial port (but don't turn it on just yet), and enter the following at a text prompt **pilot-xfer -i newsoftware.prc**. Of course, change **newsoftware.prc** to the proper name of the program you're trying to install.

Depending on your setup, you may also need to enter the above command as the root user. You'll see a message directing you to press the HotSync button on your Palm, and then the transfer will take place. After it has finished, you're returned to the prompt where you can safely disconnect or upload something else.

This is the simplest use of the pilot-tools package, and is a fast and easy way to transfer information around. You can install any program you download with **pilot-xfer -i <name.prc>** but this is only a small part of it. Check out the documentation to see what else pilot-tools can do.

4. WORKING WITH JPILLOT

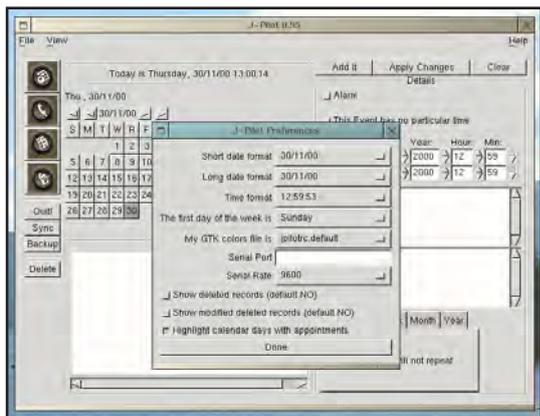
Now that you're a little familiar with the fundamental *pilot-link* programs, you can try out one of the graphical applications which emulate the *Palm Desktop* software supplied for Windows machines. One of the best is *JPilot*, freely available under the GPL and downloadable from the project's website: <http://jpilot.linuxave.net/>

Here you'll find RPM packages and the source, and the only major requirements are the Gtk+ toolkit libraries along with the pilot-link suite of tools. If you're building from source, you'll also need the pilot-link-dev package too.

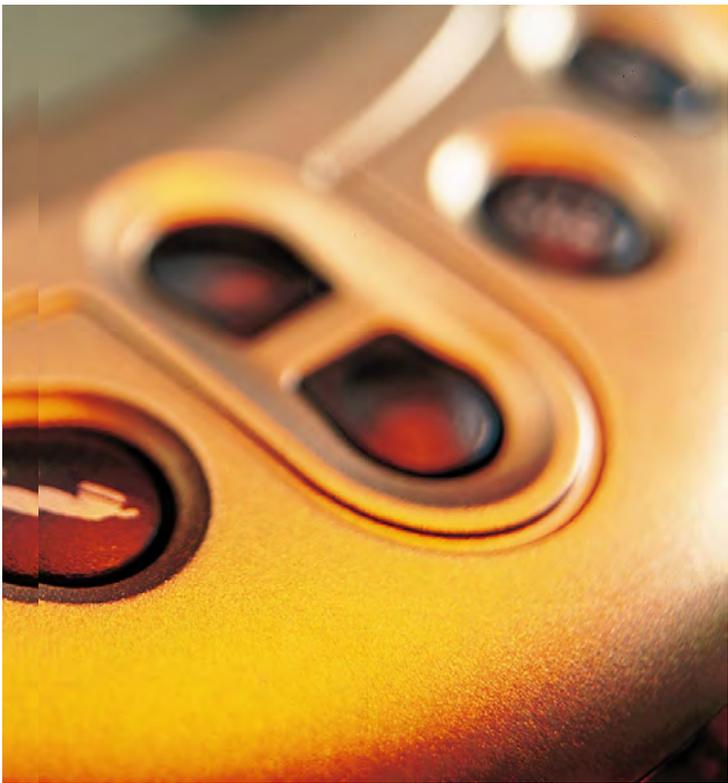
JPilot doesn't boast all the features that its Windows cousin provides, but it does an excellent job of synchronising your Address Book, Schedule, Memos and To-do List. There's also an email plugin if you want to compose messages away

The pilot-link tools in action at the console, with an application and text file being installed.

```
nike@localhost nike$ ls /usr/bin/pilot*
usr/bin/pilot*      /usr/bin/pilot-debug*  /usr/bin/pilot-undelete*
usr/bin/pilot-addresses*  /usr/bin/pilot-file*  /usr/bin/pilot-xfer*
usr/bin/pilot-clip*      /usr/bin/pilot-mail*  /usr/bin/pilotListener*
usr/bin/pilot-debug*    /usr/bin/pilot-schlep*
nike@localhost nike$ pilot-xfer -i ykConc.prc
aiting for connection on /dev/pilot (press the HotSync button now)...
connected
nstalling ykConc.prc... OK
nstall done
nike@localhost nike$ install-memo up_roundup.txt
lease Insert Pilot in cradle on /dev/pilot and press HotSync button.
nike@localhost nike$
```



JPilot busy at work, showing the schedule window and preferences box.



from the desk. Of course, it sports the usual file-installation and removal as well, and on the whole is very straightforward to use.

If you fire up the program and look down the left hand side of the main window, you'll see four buttons with an appearance similar to those on the Palm device itself. These correspond to the main PalmOS organiser utilities, and clicking on each button brings up a familiar layout for that specific function.

So, you can hit the "Sync" button to fetch the data from your PalmOS machine and start editing it in *JPilot*. When you've finished, simply hit Sync again and the records on your handheld will be updated. Easy!

To enter information in the Schedule/Datebook, navigate through the calendar using the panel in the top-left. You can click on individual days, and when you need to input an appointment, move over to the right-hand Details pane and select the correct date and time. You can also set up an alarm for this event too.

The Address Book is equally easy to operate, and you can filter out certain types of messages (Business, Personal etc.) and add new entries with the pane on the right. This has three tabs for the various details you can enter – also, the drop-down list along the top helps you to choose the appropriate category.

JPilot's Memo section uses a basic text editor to input new messages – the ability to paste text from other

applications is a nice touch too. On the File menu is an entry for installing new applications, together with the Preferences box which lets you modify the time/date formats and so on.

5. OTHER ALTERNATIVES

As you might expect, *JPilot* isn't the only Linux tool for working with PalmOS devices. If you're a KDE user, an application that's worth considering is *KPilot*. This Qt-based program provides similar functionality to *JPilot*, and is equally user-friendly. It may well be sitting somewhere on your distro CDs – if not, take a trip to the project's website at: http://www.slac.com/pilone/kpilot_home/

KPilot's drop-down list on the right gives you access to two of the main Palm organiser functions, and on the left is the HotSync button (for when you want to transfer the data across). You can import Memos and even Address Book lists from text files on your hard drive.

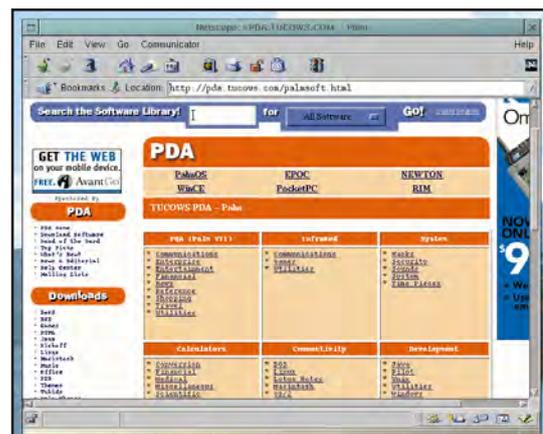
If you check out the "Setup" entry in the "Conduits" menu, you'll see that a number of additions can be installed which extend the capabilities of *KPilot*. The "PopMail" conduit, for example, lets you use the handheld for composing and reading email (with options to modify the SMTP/POP servers, signature file and more).

6. HAVE FUN!

We've gone through the various stages of installation and setup of your PalmOS communication software. After following this tutorial, you can now make full use of your Linux system and Palm handheld together – and you no longer have to rely on Windows to keep your pocket pal up-to-date.

We've looked at two of the main PalmOS desktop applications here, but there are plenty more in development and you may come across something else that takes your fancy – have a glance at the Linux and PalmOS on the Net box for a few pointers, and if you find something really interesting, let us know. Happy hunting! 🐱

KPilot is a KDE tool for communicating with your Palm handheld.



Check out Tucows' PalmOS pages for a wealth of top programs.

LINUX AND THE PALM OS ON THE 'NET

There's a whole stack of Linux-related info on the Internet for those using a Palm handheld. You'll find documentation galore for developers, users and those just looking to grab the latest software. Your first port of call for all things Palm related should be the Tucows PDA pages at: <http://pda.tucows.com/palmsoft.html>, which houses more development software and utilities than you'll ever need. If you fancy doing a spot of programming for the PalmOS, you'll be chuffed to find a number

of tools and guides which explain the various requirements and details needed to write your own applications.

However, if you're more interested in using than creating, the Tucows site features a very large selection of the best PalmOS software available today, and once you've got the pilot-link package up and running, you can dive straight into the masses of games, utilities and full-blown applications waiting patiently to be downloaded onto your handheld.

NETWORK FILE

Jon Kent looks at the configuration and use of the Network File System (NFS), and how it can provide simple, if insecure, filesharing.

Linux Info

Most of NFS is usually included as part of the base install. The server and filesystem are kernel modules, so don't worry about downloading it from anywhere – usually all you'll have to do is turn the service on!

NFS was designed by Sun Microsystems in the early 1980s and was standardised in 1987 (RFC 1094). As NFS is an open standard it is not just limited to UNIX systems but can also be run on multiple platforms. Simply put, NFS allows you to access remote directories and files as if they were located in your local machine. You can perform all functions (add, remove, edit and so on) on files that are NFS mounted, assuming that you have the required permissions. NFS is commonly used to supply user's home directories and access to applications. The abilities that the system provides to access remote drives has many advantages for the users and administrators of a GNU/Linux network:

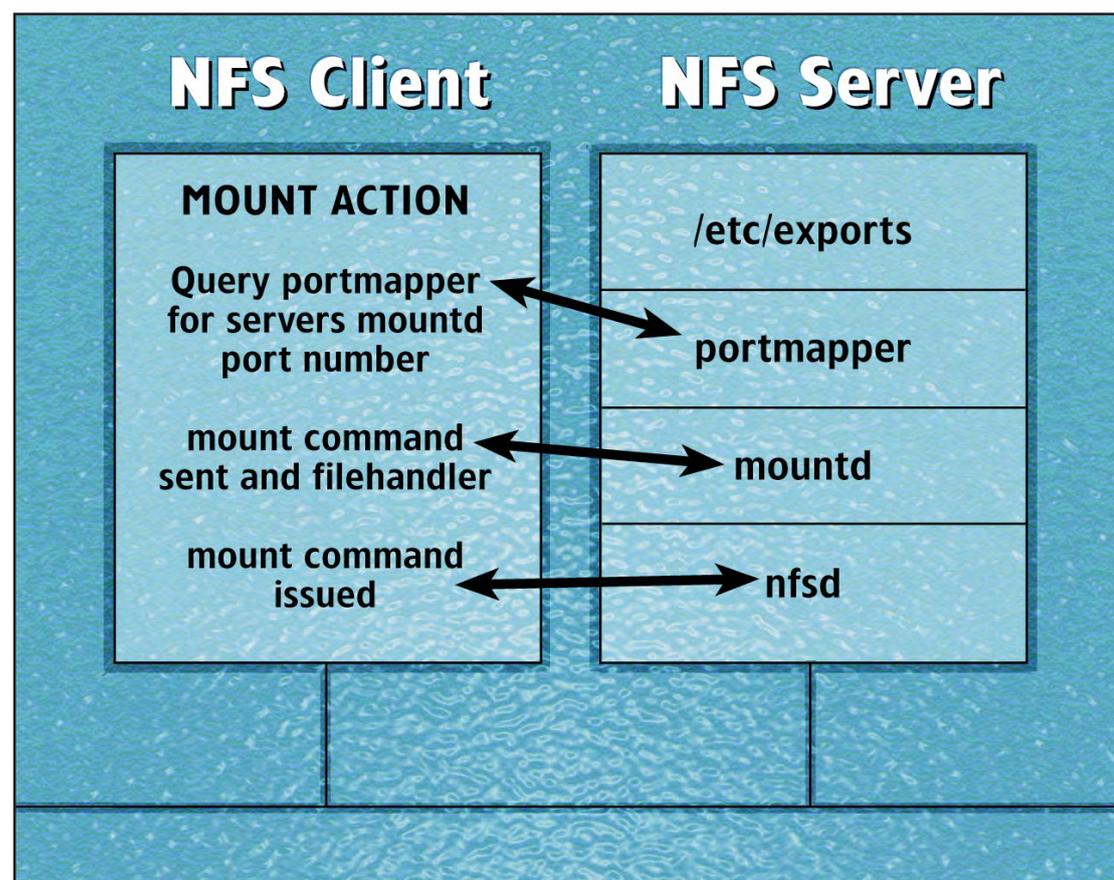
1) Simpler management of applications, as application can be installed on a NFS mount point for the users to access, and do not have to be installed locally on all interested users workstations.

2) A centralised backup and restore system, as all user home accounts, application and other such data can be held on a server or servers controlled by the administrator

UNDER THE HOOD OF NFS

NFS uses Remote Procedure Call (RPC), is stateless, communicates over User Datagram Protocol (UDP) – though some non-GNU/Linux versions can also use TCP – and is based around a client-server architecture. RPC provides a framework for a user's local application to call a set of procedures on a remote server. When an RPC function is called by a local process (client), the function is packaged up and sent over the network to the remote server in question. The remote server will then run the function that the client program has requested and return the results back over the network to the client.

RPC does not use 'well known ports', which are controlled



The relations between the NFS server and the client, and what file controls what!

SYSTEMS

by inetd, so a separate daemon from inetd is used called portmapper (rpc.portmap). The portmapper is an RPC server and is assigned a 'well known port' which is port number 111. RPC programs register the port they intend to use with the portmapper, which in turn allows applications that wish to use the RPC services to query the portmapper for the port number of the RPC service. Although NFS could in theory use any port via the portmapper, it usually uses 2049.

NFS is referred to as a stateless protocol, that is, no client or server can go into a state that depends upon further information. This is possible because all information required is provided within the functions parameters that are passed to the server and back to the client. Therefore no external information is required for the function to complete. This statelessness adds a degree of reliability to NFS.

As mentioned NFS (and RPC) is used by a number of platforms. In order to support these differing platform XDR (External Data Representation) was developed. Like NFS, XDR was developed by Sun Microsystems and is an open standard (RFC 1014). XDR defines a framework that must be used to encode values in an RPC message. Because of this standardised framework, passing RPC messages between differing platforms is possible.

In GNU/Linux, NFS uses UDP (User Datagram Protocol)

as the transport mechanism, encapsulated with an IP packet. UDP has no method to guarantee data delivery, which may seem an odd transport mechanism to use, however, using UDP provides several advantages because of its connectionless nature. The primary advantage is that NFS can

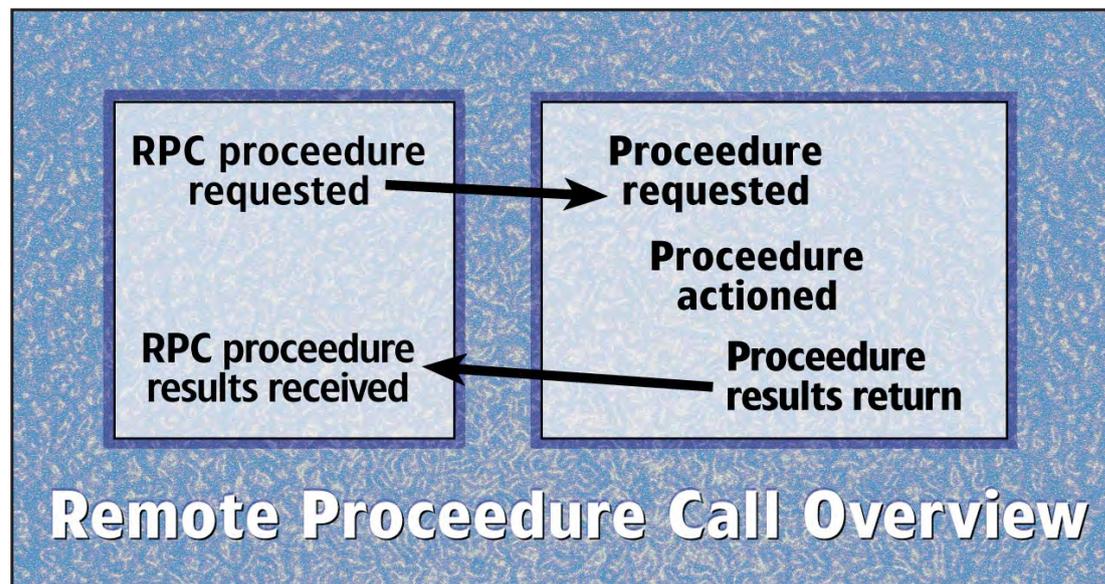
survive network failures, or the temporary loss of the server from which it was being mounted.

In addition to the above, several other RPC programs are part of an NFS system. These are the NFS daemon itself (nfsd or rpc.nfsd) and the mount manager (mountd or rpc.mountd). On an NFS server, the daemons nfsd, portmapper and mountd are used

to serve (or export) some of its filesystems to be NFS mounted by clients. An NFS client also uses the daemons mountd and portmapper to mount NFS exported filesystems. In addition to these daemons an NFS client can also run another program called automounter (or autofs) which is explained later.

NFS support under GNU/Linux is configured in the kernel. There are two options that can be set within the Filesystems screen in the Kernel Configuration program, these being Kernel automounter support (explained later) and NFS filesystem support. Kernel automounter may need to be configured, depending upon your requirements, either as a module or directly into the kernel itself. NFS filesystem support is required, again as a module or directory in the →

NFS is a stateless protocol: no client or server can go into a state that depends on further information



Web resources

You'll find plenty of NFS info on the web with a simple search, but a good place to start in <http://nfs.sourceforge.net>. As well as the latest versions, you'll also find a comprehensive FAQ and user guide.

Passing RPC messages between different platform is possible thanks to the XDR protocol.

TUTORIALS | network file system

→ kernel. Strictly speaking this does not have to be configured within an NFS server's kernel, but there is little overhead in this. Use a kernel module if you trying to lower the kernel memory footprint.

CONFIGURING AN NFS SERVER

Configuring a server to act as an NFS server is straight forward. First, you need to define which directories you wish to export (made available to be mounted), which are defined in the file `/etc/exports`. The basic syntax of this file is:

```
/home (rw)
/apps linux*.my.domain (rw) clay (ro)
/phone (ro)
```

The `exports` file has many possible options that can applied to an exported filesystem, refer to **man exports** for more detail. This example is broken down in the boxout below:

Once the `exports` file has been created you can then

either reboot the server or run the relevant startup script located in `/etc/init.d`. On a Debian system the startup script is called `nfs-kernel-server`, however, the name of the startup script (and

possibly the directory it is in) may be different in other distributions. On a Debian system you would run the following to startup the NFS service:

```
/etc/init.d/nfs-kernel-server start
```

Once the relevant script (or reboot) has been actioned, the daemons `nfsd` and `mountd` should be running and the filesystems defined in `/etc/exports` should now be exported.

If the `exports` file is modified after the initial installation, you can either restart the NFS server and **mountd** to export the new entries or use the command **exportfs**, if available. To use the command `exportfs`, you simply have to update the `exports` file and then run:

```
exportfs -a
```

This will then update the NFS server with the additional exports that have been configured.

The file `/var/lib/nfs/xtab` is created, or updated, when a

new export is created. This file is then read by `mountd` when a remote host requests access to filesystem to check it's available.

OPTIONS AND EXPLANATIONS

<code>/home</code>	<code>(rw)</code>	export <code>/home</code> directory as a read-write export, any host can mount this export (NOTE this is insecure).
<code>/apps</code>	<code>linux*.my.domain (rw)</code>	allow only hosts that match the pattern to mount this export.
<code>/apps</code>	<code>clay (ro)</code>	only allow host <code>clay</code> to mount this export as read-only.
<code>/phone</code>	<code>(ro)</code>	export <code>/home</code> directory as a read-only export, any host can mount this export.

CONFIGURING AN NFS CLIENT

There are several ways to mount a remotely exported directory onto a local machine, these being manual mounting, an entry in `/etc/fstab` file or automatic mounting using automounter. To manually mount an exported filesystem you would enter the following command:

```
mount 172.18.42.120:/opt/gcc /gcc
```

This would mount the exported directory `/opt/gcc` from the server (in this case denoted by an IP address, `172.18.42.120`) onto the directory `/gcc` on the local machine. The directory `/gcc` does have to exist locally for the mount command to work. To unmount this NFS mount you run the command **umount [mount point]**.

To have an NFS export mounted every time you reboot your workstation you need to put an entry in the `/etc/fstab` file, for example:

```
server1:/opt/gcc /gcc
```

This entry performs the same function as before in the manual mount, but will now mount this export upon a

reboot. There are several options to the mount command which you may wish to use, refer to the man page for more details.

Lastly, kernel automounter (`autofs`) support allows filesystems to be automatically mounted as they are required. Automounter uses automounter maps that can either be held locally on the workstation (usually in `/etc`) or forming one (or more) of the tables within an NIS (Network Information Service) environment. As mentioned previously, automounter support is provided by the kernel and therefore you need to ensure that you kernel is compiled with this option. There can be several configuration files to control automounter, however there is a master configuration file called `/etc/auto.master`. The `auto.master` defines the local mount point to be used and the corresponding automounter map that points to the export(s) that can be accessed. An example of an `auto.master` configuration file using local maps, not NIS maps is:

```
#
# Sample /etc/auto.master file using local maps
#
/opt /etc/auto.opt
/home /etc/auto.home
```

If the automounter maps are available via NIS, the `auto.master` file would look like this:

```
#
# Sample /etc/auto.master file using NIS
#
/opt auto.opt
/home auto.home
```

With the above configuration all access to the these mount points will lead to the associated maps to be referenced to verify where the required export is located. Therefore, if a user tries to access `/opt/gcc`, the automounter map `auto.map` will be referenced to locate the export that is to be used for `gcc`.

The automounter map file uses the following format:

```
[export] [option(s)] [NFS server:]/[path]
```

For example:

```
#
# Sample auto.opt
#
gcc -defaults,soft,intr server1:/opt/gcc
bin -defaults,soft,intr server2:/opt/bin
```

Automounter is the cleanest method for users to access NFS directories, as it means users do not have to worry about which server an export is accessible from, and how to use the `mount` command. If used in conjunction with NIS or a regular **rdist** or **rsync**, automounting can make administering an NFS environment easier to control.

SECURITY

As with any network based service there are security issues that need to be taken into account. To understand some of these issues, the commands **rpcinfo** and **showmount** need some explanation.

As NFS is an RPC based service, the portmapper can be remotely probed to discover which rpc services are being run using `rpcinfo`. Running `rpcinfo` as follows would display the rpc services running on a remote server:

```
rpcinfo -p [server]
```

This will display a list of programs and their assigned port number. If this is run against a Linux box running as an NFS server, the program `nfs` will appear in this listing in addition to `mountd`. Now, to display the exported directories from this server, you run the **showmount** command as follows:

```
showmount -e [server]
```

This will now display a list of exported directories, and the mount permissions set on these exports. Should a directory be exported without any mount permissions this directory could be mounted onto any machine that can access the server. The point of this example is to show how easy it is to discover which servers are NFS servers, and to query the exports of these servers. Therefore it is extremely important not to export a directory without any mount permissions set.

NFS Protocol Family Tier

NFS

XDR

RPC

UDP

IP

Additionally NFS should never be run in a secure environment, such as a DMZ.

There are a few extra export options which can be set to reduce the risk of a server being attacked. The first and probably most important one is **root_squash**. As far as NFS is concerned a user is a user, so if you are logged in as root on a remote machine, you will have root access to the exported directory on the server too, not particularly ideal. The following exports line will address the problem:

```
/home (rw, root_squash)
```

The NFS server will now automatically map any remote root user to the anonymous user 'nobody'. The idea can be extended further by using the command **all_squash**. This maps all remote users to the anonymous user, or to a specific user if you prefer:

```
/home (rw, all_squash, anonuid=540,anongid=100)
```

This maps all remote access to the user and group ids specified. Further security options are explained in the exports man page.

In conclusion, NFS provides a method to centrally manage common directories, such as `/home`, and applications. It is simple to install and manage, but does require some thought to any security issues that may be raised. 

This diagram shows the layers of protocols used when running an NFS server.



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

ANSWERS

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

Twice the fun?

Q I am considering the purchase of a reconditioned Compaq w5100 and using its dual processor capability, with a view of installing Linux as a server for a small network of eight computers. Is this a realistic choice or should I look at a more modern machine? Also you comment on the "fairly infamous instabilities of NFS", can you expand on this?

A The specification of a server does, of course, depend upon what it is you are serving. If your just going to run a NFS server, to share home directories or common files, then you don't need much CPU power at all, just a lot of bandwidth and fast drives. I'm not familiar with the specification of the Compaq machine, but you could build yourself a high end Athlon based file server with 256Mb of RAM and over 30Gb of disc space for around £400, if that. You don't really need that much hardware, but chances are you'll end up running other network portable services on it, such as mail or a web cache, which do require a bit more kick. Second hand machines are okay, up to a point, but hardware – particularly end of line – is so cheap now, there is little reason to buy something which will have compatibility issues in the near future. What if you wanted to upgrade the RAM or replace a burnt out CPU? Sure, you could get the bits eventually, but very rarely does it warrant the amount of time, hassle and expense.

Traditionally NFS has been rather

Your kppp settings will vary depending on which ISP has the benefit of your custom.

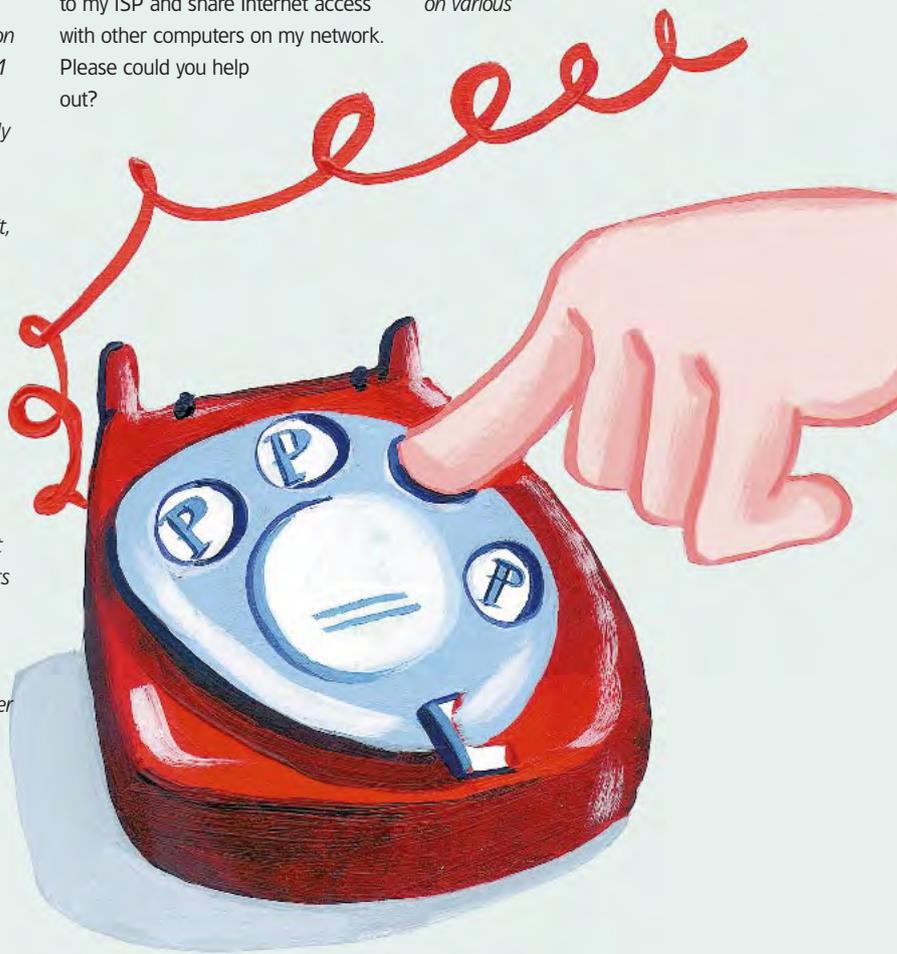
dodgy, particularly on early Linux kernels. There are a number of patches, available from <http://nfs.sourceforge.net/> for the current 2.2 kernels which provide updates NFS code and various fixes for things broken in the main kernel release. Another option, if you feel confident enough, would be to run the latest 2.4.0-test kernel on the file server, as 2.4 has the updated NFSv3 code built in, which is far better than that in 2.2. As the major stability issues with NFS are with the server, you don't need to run 2.4.0-test kernels on any other machines.

Net gateway

Q I have installed Mandrake-Linux 7.1 on my PC. I want to connect to my ISP and share Internet access with other computers on my network. Please could you help out?

A Without knowing which ISP you are using, it's difficult to give specific details, but a tool such as kppp or gnome-ppp can be used to configure a dial-up interface. Simply pop in the number to dial, username, password and any specific authentication options, and hopefully you should find yourself connected to the Internet. If your using an ISDN TA, rather than a standard serial modem, you should look in the direction of the ISDN4Linux software, at <http://www.isdn4linux.de/>.

To share the Internet connection with other machines on the network requires you to compile IP Masquerading support into the kernel. If you are unfamiliar with recompiling the kernel, there are HOWTOs on various



web sites, including linux.com/howtos/, as well as it having been discussed in earlier issues of LXF. Once you've rebuilt the kernel and rebooted, ensure that you have `ipchains` release 1.3.9 installed. If not, head along to your nearest Mandrake mirror and install it.

Assuming your LAN is using 10.1.1.0/255.255.255.0 as the network, the LAN is on `eth0` and your machine is 10.1.1.1, you simply do:

```
ipchains -P forward DENY
ipchains -A forward -i eth0 -s
10.1.1.0/24 -j MASQ
echo 1 >
/proc/sys/net/ipv4/ip_forward
```

After setting the gateway for the other machines to 10.1.1.1, they should be able to access the Internet once you have supplied them with appropriate name server IPs.

Sound advice

Q I purchased and installed SuSE Linux Personal Edition. The installation was easy and I very much like what I see. The problems I have is that Linux knows I have a SB128 sound card but says it can't play sounds, and refers me to check if it's ALSA compatible or something sounding like that.

Also, Linux does not recognise my modem (a Supra Express 56i pro) at all.

PS I really like the look of Linux and can see it taking off soon in a big way.

A The sound card is fairly difficult to install as you need to download kernel modules from Creative Labs' site and install them by hand. Everything you need to know, including where to get them from and how to install them is at <http://www.linuxnewbie.org/nhf/intel/so undcards/sblive.html>

You'll not have as much luck with the modem, as it is a CPU-powered WinModem. If you're desperate, feel free to try the kernel modules at www.linmodems.org, but your mileage will vary and if you intend to use Linux



GNOME or KDE? Both Linux desktop environments have their dedicated followers.

to access the Internet, the purchase of a external serial modem would be in your favour.

As for the last comment... I think you'll find that there is the odd company here and there, trying to make a buck or two out of Linux. Suffice to say, from a server point of view, it certainly has already.

Window swapping

Q I have a problem with KDE and the desktop windows. Whenever I go near the edge of the screen, I seem to end up being flipped or changed into a completely new desktop. Is there some way to stop this?

A What's happening is that you have active desktop borders enabled for KDE. This means that KDE sets multiple 'virtual' desktops, and allows you to switch between them by moving the mouse 'off' the edge of the screen.

You can adjust the sensitivity of this effect, or turn it off in the KDE Control Centre, under `LookNFeel/Borders`. If you wanted to keep the effect, but make sure you really wanted to switch screens, the easiest thing to do is increase the delay, so you will have to purposefully move your mouse to the edge of the



screen and keep it there before the switch occurs.

Big drives

Q I have a 2Gb Seagate drive which I've just installed in an old tower (486), hopefully to use as some form of server. The problem is that this old piece of junk has a 520Mb drive limit. Can *Lilo* handle this problem like the Seagate boot manager, or do I have to have that running, then boot *Lilo* to start Linux?

A The drive limit is probably a BIOS restriction - can you update your BIOS? This won't effect *Lilo* or *grub* or whatever else you want to install.

Triple Boot

Q Most people have a dual boot system, but I want to set up a triple-boot system with Windows 98, Windows NT and Linux. I have plenty of drive space, but I'm wondering how I'll be able to boot into all three systems. My mates have told me it is possible, but I'll have to boot Linux from a floppy.

A Not necessarily so. The easiest way around this is to install Windows first. If you subsequently install Linux, you can use *Lilo* to either

boot from NT or Windows 95. *Lilo* should detect the presence of both OSs, or you can manually edit *Lilo* and add an entry for NT, such as :

```
# Windows NT boot
other=/dev/hdc1
label=/dev/hda
loader=/boot/any.db
label=Windows NT
# End Windows NT
```

Taking a Gnap

Q My version of *Gnapster* doesn't seem to be working anymore. It used to be fine, but now it just sits there forever, saying that it's trying to find the best server, before returning an "unknown fatal error"

A Hmm, you don't say what version of the Linux client you are using, but it is possible that it's out of date. *Napster* changed some parts of it's protocol, so if your version is below 1.3.12, it probably won't work. Get the latest version from jasta.gotlinux.org/gnapster.html

X crash

Q I run Mandrake 7.1 and tend to try a lot of early release software from sourceforge just for the sake of experimenting. As to be expected a lot of this is not so stable and I get a fair few crashes. I'm sure I read somewhere that I probably just crashed X and not Linux itself. If that's true, then how do I restart X if the machine is locked? The mouse still moves but that's all.

A If the keyboard is still being scanned, then you can kill X with `Control-Alt` and the Backspace key. This may or may not work, depending what has gone wrong. If it doesn't, try `Control-Alt F2`. This should bring up a terminal, where you can log in, find out what process number X is by typing: `ps x` and then kill it with: `kill <uid>` where <uid> is the process number you just found.

If that doesn't work either, then the keyboard isn't being scanned, →

→ and the only realistic way in is to telnet in to the box from another machine on the same network, if you have one.

Do yourself a favour and install the ReiserFS, so hard resets of the box don't cause so many disk problems...

Canon Can

Q Anyone know if it's possible to run a Canon BJC 3000 printer with USB connection in Linux Mandrake 7.2. At present my family use Windows Me with the printer set up in USB. If I can't run USB in Linux, can I use the spare printer port on the back of the Canon? Can I use this port without removing the USB connection in Windows. This would stop my family screaming at me for messing up the printer again.

A You should find that this printer works fine on the USB port under Mandrake 7.2. At least, Mandrake will find the device. Whether or not you can print to it will depend on the driver. Mandrake 7.2 also includes the CUPS system, and Gimp-Print, which provides CUPS compatible drivers for a whole range of printers (both of which are also on our CD this issue).

The Gimp-print driver for the BJC3000 is believed to work, but is untested at the moment, so I'm sure if you get it working, the developers would appreciate if you surfed on over to gimp-print.sourceforge.net and let them know.

If, for some reason, Mandrake can't talk to your printer over USB, it should be fine to have it connected via the parallel port. Most printers are happy with this, as long as you don't try communicating by both interfaces simultaneously!

Reader Response

Q I have enclosed some details that may help Kim Beadle (and maybe some others) sort out the Modem set up, perhaps you would forward it on.

I have configured a modem (US Robotics 56K voice modem 9190 & a Creative labs sound card (Vibra 16) with both Red Hat and Mandrake. Both were quite difficult, but achievable. Configure the sound card first using *Drakconf* or the command **pnpdump -c** from a command line. This creates file in /etc called *isapnp.conf*. This has the settings for the card and the modem. Alternatively, use *Drakconf* in Mandrake and adjust the settings for the DMA8 and DMA16 until you get a sound test. The same file is created using this method. My settings are:

```
(CONFIGURE CTLO070/-1 (LD 0
(DMA 0 (CHANNEL 3)) (DMA 1
(CHANNEL 7)) (INT 0 (IRQ 5 (MODE
+E))) (IO 0 (SIZE 8) (BASE 0x220))
(NAME "Creative Labs SoundBlaster
Vibra 16") (ACT Y) ))
```

Set up the modem using Windows, if you have it, and read the modem interrupts and IRQ from the My Computer>Properties. Then boot into Linux.

To set up the modem, copy the **/usr/doc/setserial 2.X/RC.serial** file to /etc/rc.d/init.d/serial (read setserial manual pages (setserial man) and Setserial HOW TO),

create a file called **serial.conf** in /etc, with a text editor type at the first line of "serial.conf"

```
###AUTOSAVE###
```

At a command line type **Setserial ttyS0** then repeat for ttyS1, ttyS2, ttyS3, taking note of the settings.

I set my modem to run to ttyS2, IRQ 3, UART 16550a 0x03e8 ensuring that there are no conflicts (you may need to play with these to get the modem to work after you have changed the linuxconf settings).

Next run Linuxconf and change the modem setting to **ttysX** i.e the tty (COM port) you have configured your modem to e.g. **ttys2**. At a command line type **/etc/rc.d/init.d/serial stop**. This will save the setting to a file called *serial.conf*

Test the modem, and if it says 'modem busy' try changing the **0x0** setting or IRQ or ttySX. If it says 'modem not found', it is possibly the incorrect ttySX is symlinked. After each change of the serial setting repeat **/etc/rc.d/init.d/serial stop**.

If the modem works, copy the files *serial.conf*,

isapnp.conf and *isapnp.gone* to a floppy for safekeeping in case you reinstall the system.

As for the ISP information, you will need to enter the Ip address of the ISP into KPPP (Freeserve & BT are quite good at providing this information - why not change)

Reboot the system.

Hopefully this will result in the soundcard and modem working.

S J Walker

Flashing lizard

Q I have just installed Mozilla, but I'm having trouble viewing some sites that previously worked fine under Netscape. Flash doesn't seem to work at all. Are there any tips you could give me?

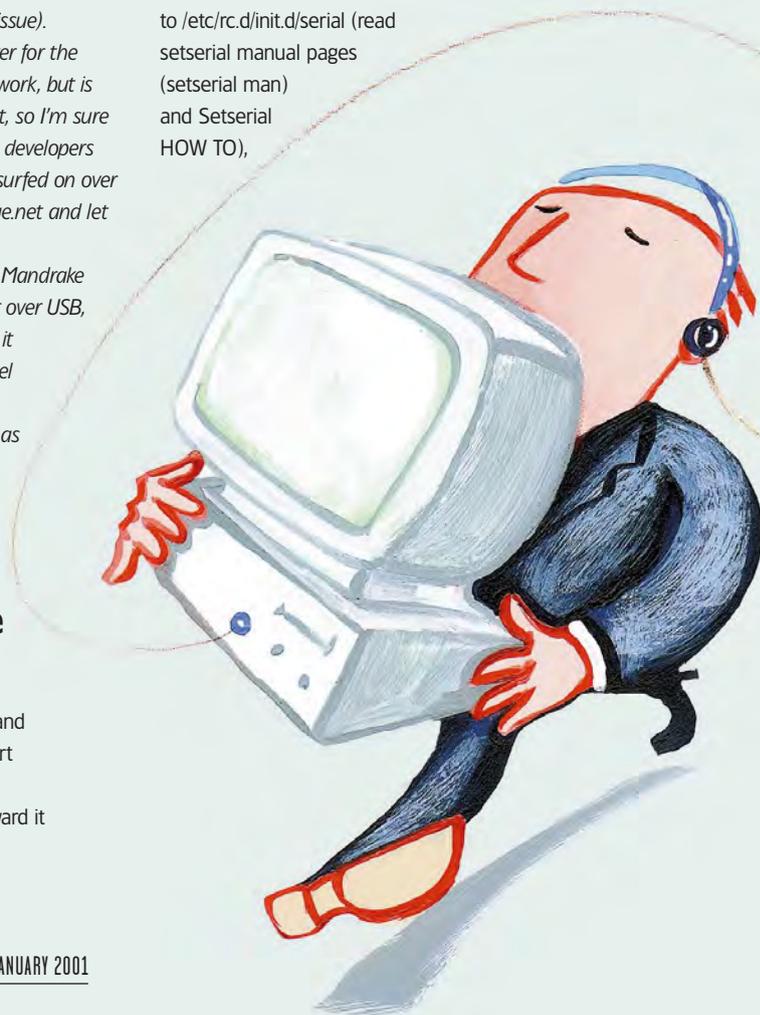
Q You can use pretty much any of the Netscape plugins with Mozilla. In fact, you're probably best off copying all the plugins from the *netscape/plugins/* folder to your *Mozilla/plugins/* folder and ditching Netscape altogether!

Alternatively, download the plugin from the Metacreation site (www.metacreation.com), uncompress it (for some reason it is labelled *flash.tar.gz*, but it isn't actually zipped) and copy the Java class and library into the Mozilla plugins folder. It should work fine. You can check it is installed by typing *about:plugins* as the URL

More modem problems

Q First I'd like to thank my best Linux magazine. I have read most issues of your LXF magazine. But I couldn't find the issue that had SuSE 7.0. I have P233MHZ, 64Mb, ESS1868 sound card, S3 virge 4MB, Rockwell HCF 56 PCI modem (on COM3); I installed WinLinux2000 over Windows ME.

1. WinLinux2000 has detected my modem, but it still not working properly. I get a message 'modem is busy' under kppp dial up, but when I try to correct this problem by using minicom and change tty1 to tty2, my computer hung. I'm a new comer to





Linux so be as detailed as you possibly can.

2. WinLinux2000 does not detect my sound card, so I forced it with the same IRQ, DMA, IO, and Memory range that Windows used but it still not working. When I try to use mixer I get a message 'do chmod a+rw /dev/mixer' and I do it properly, but nothing happens.

3. When I try to update WinLinux2000 kernel 2.2.13 with the version that comes on your CD, by using 'patch-linux' I get a message 'cannot detect kernel version'.

4. Is there an Arabic version for Linux? I found StarOffice from Sun better than MS-Office, however it doesn't support Arabic. Is there a version that does?

Abbas Al-Mutawa,
Bahrain

Thanks for your letter, Abbas. Unfortunately, we don't have any copies of issue 6 left, but you can get generally buy back issues from our Subscriptions Department. Contact them by email at subs@futurenet.co.uk or phone at +44 (0)1458 271178. Now to your questions.

Firstly, though, nobody at Linux Format is particularly familiar with Win4Linux, so while the answers below should apply to any version of Linux, we don't know precisely how Win4Linux differs from a standard distro. You should check out any

documentation supplied with your distro.

1. Your modem is a so-called 'Winmodem' and is therefore not supported by Linux at all. You should consider buying an external modem instead, since any external model will work happily with Linux.

2. Your soundcard, however, is definitely supported under Linux. The ESS1868 is an ISA plug-and-play device and will work with the kernel's Soundblaster driver. Configuration may be tricky, though. You need first to set up Linux's plug-and-play tool to initialise your card. This can be done by adding an entry like the following. (This assume that the main card I/O is at 0x0220, the FM synthesiser at 0x0388, the MIDI port is at 0x0330 and the card uses IRQ 5. You may have to change these to suit your set-up.)

```
(CONFIGURE ESS1868/-1 (LD 1
(IO 0 (BASE 0x0220))
(IO 1 (BASE 0x0388))
(IO 2 (BASE 0x0330))
(DMA 0 (CHANNEL 1))
(INT 0 (IRQ 5 (MODE +E)))
(ACT Y)
))
```

to the file /etc/pnp.conf. This file can be generated with pnpdump command.

See the Plug-and-play HowTo www.linuxdoc.org for more information about Plug-and-play configuration with Linux.

You then need to tell your machine to load and configure the relevant kernel modules to use your particular card.

```
/sbin/modprobe sound
/sbin/insmod uart401
/sbin/insmod sb io=0x220 irq=5
dma=1 dma16=-1
/sbin/insmod mpu401 io=0x330
/sbin/insmod opl3 io=0x388
/sbin/insmod v_midi
```

You may find it easier just to switch to a more modern Linux distribution, such as Mandrake 7.2 or SuSE 7.0, which does a better job of configuring audio for you.

3. The kernel source that we have provided on past coverdiscs have been complete source archives, not patch files. To build a kernel image, unpack the archive, and CD into that directory. Configure the image to build by executing /make menuconfig/ or if you are running on X, /make xconfig/. Then enter the series of commands

```
make dep
make bzimage
make install
make modules
make modules_install
```

to build and install the new kernel image. You should really know what you are doing before installing a new kernel image, however, so either see our Kernel tutorial in issue 3 or check out the Kernel HowTo.

4. You are probably better-placed for finding out about Arabic versions of Linux, since there is very little information available on this topic which is not written in Arabic. You could try the web-page www.linux4arab.com. It looks interesting, but obviously we can't read it! Debian Linux (www.debian.org) also have an Arabic page. As for a

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 baked beans 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

word processor, we don't know of any available yet. The KDE project have a translation team working on an Arabic localisation, but I suspect it will be some time before a translation of KWord is ready to go. Can any of our other readers who speak Arabic help Abbas here? **LXF**

BENEFITS – WHY SUBSCRIBE?

- ✓ Save over **£1** on every issue*.
- ✓ Keep your money in your pocket longer when you subscribe by Direct Debit. It's easy! Complete the form below or call 01458 271178 and have your bank details ready.
- ✓ Delivery direct to your door.
- ✓ Never miss an issue.
- ✓ All you need to do is fill in the form. If you live in the UK, you don't even need a stamp!

RISK-FREE GUARANTEE
You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked.

Title _____ Initials _____ Surname _____
Address _____
Postcode _____ Country _____
Email address _____ Daytime telephone number _____

Please choose your method of payment **1, or 2**

- 1.** Direct Debit (UK only) – I wish to pay £24.95 every 6 months.
I understand that I will receive 13 issues during the next 12 months.

Instruction to your Bank or Building Society to pay Direct Debits.



Originator's Identification Number

7 6 8 1 9 5

Future Publishing Ltd, Cary Court,
Somerton, Somerset, TA11 6BR

Please send this form to address shown

1. Name and full postal address of your Bank or Building Society branch

To: The Manager _____ Bank/Building Society
Address _____
Postcode _____

2. Name(s) of account holder(s)

3. Branch sort code
(from the top right-hand corner of your cheque)

4. Bank or Building Society account number

5. Instruction to your Bank or Building Society

Please pay Future Publishing Direct Debits from the account detailed on this Instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s) _____ Date _____

Ref No (Office use only) _____

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

LINUX FORMAT

The UK's **FIRST** dedicated Linux magazine

2. Cheque / credit card 13 issues

UK	£49.90	save £14.97
Europe (including Eire)	£73.53	
North America	£73.53	
Rest of the World	£86.53	

Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)

Visa

MasterCard

Switch

Issue No/Valid date _____

Card No _____ Expires _____

Signature _____ Date _____

UK readers return this coupon to:

LINUX Format, Subscriptions, Future Publishing Ltd,
Freepost BS4900, Somerton, Somerset, TA11 6BR.

Overseas readers return (postage payable) to:

LINUX Format, Subscriptions, Future Publishing Ltd,
Cary Court, Somerton, Somerset, TA11 6TB, UK.

Internet address:

www.futurenet.com/promotion/pg010/92

Please tick this box if you would prefer not to receive information on other offers.

ORDER HOTLINE ☎ 01458 271178

QUOTE ORDER No. **LXFP10**

OFFER ENDS: 20.03.01

3 issues FREE*

*Subscribe today
to LINUX Format
and receive
13 issues for the
price of 10
-that's 3 issues free.*



**DON'T DELAY
SUBSCRIBE
NOW!**

You have a choice: you can trudge down to the newsagents every four weeks and pay £4.99 for **LINUX Format** (great value in itself) – assuming, of course, that the magazine hasn't already sold out. Alternatively, you can sit back in your comfy chair and pay as little as **£3.83** an issue, which is then delivered to your door. You'll never miss another issue and you won't give yourself a hernia getting the magazine each month!



ORDER HOTLINE: Quote order code LXFPI0 when you call

01458 271178

*UK price only

Linux Format COVERDISC

Nick Veitch introduces you to the tenth Linux Format CD - full of essential Linux software and resources...



Welcome once again to the Linux Format CD. We have packed this CD full of the latest open source projects, demos of commercial products and useful files and documents of all description. We hope you'll find it all interesting and/or useful. I'd like to remind you that we do reserve an area of the CD for Reader requests - if there is some particular project or files that you'd like to see, just let us know - see the boxout on these pages.

Omnis Studio 3.0 Lite

Want to create programs without the hassle of programming. Then give this excellent Rapid Application Development suite a try.

Omnis Studio provides an environment where you can build and test database-driven applications, simply by piecing together program components from the large toolbox provided and modifying their settings and behaviours. All this is done in a comfortable drag-and-drop, graphical interface. In many cases not a line of code has to be written, and to make things even easier, some components are supplied with wizards

IMPORTANT NOTICE

Before you even put the CD in your drive, please make sure you read, understand and agree to the following:

The Linux Format CD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of CD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.



to carry out the configuration. If you feel up to the challenge, however, you can learn *Omnis Studio's* special object-oriented scripting language and take your applications to the next level. Whichever, full documentation and some excellent tutorials are provided on the CD.

Installation

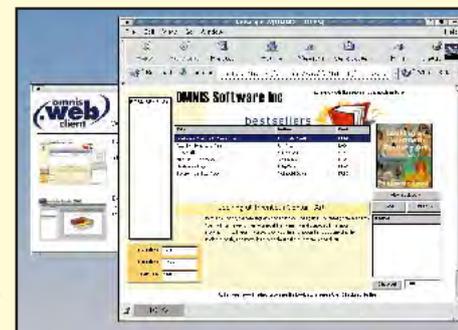
Omnis is provided on the coverdisc as an RPM package and a tarball located in the linux/install/development drawer of the Omnis distribution. If you have an RPM-based distro, then installation is simple, either

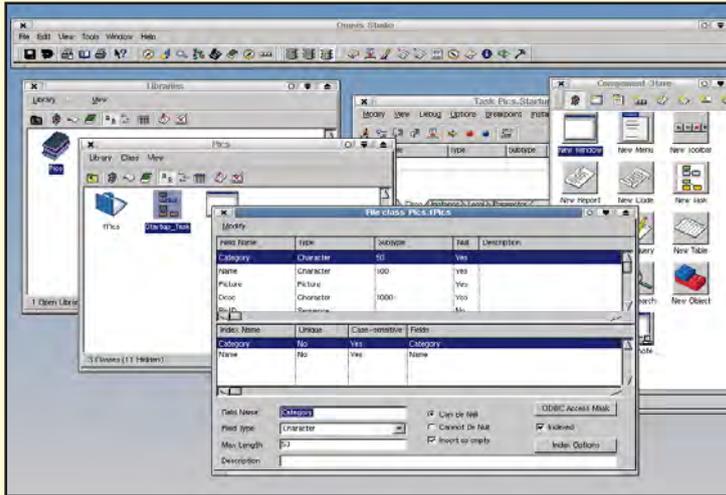
Omnis provides easy drag and drop creation of web-based applications. Try it out now.

with the installer provided or with your favourite package manager. For the former, log in as root, cd to the drawer linux/english and execute the installer with **`./setup.eng`**

For instructions on how to install from the tarball see the Readme on the coverdisc.

If you have installed from the RPM then Omnis can be started with the command **`/usr/local/os30/omnis1386`**





Omnis Studio Lite is one of the easiest rapid application development programs, and it's yours free!

Because of the way that Omnis was designed, you also need to be logged in as root to run the program. When the Welcome screen appears, click on New User for a gentle introduction to rapid application development.

Registering

The first time you start Omnis, you'll be asked for a registration code. You have two choice here: either click on OK without filling in any details - the program will start in a demonstration mode - or proceed to the Omnis website at www.omnis.net to obtain a free registration code for Omnis Lite. Enter your name, email address and "LinuxFormat" as the magazine code, and you'll be emailed a registration code to unlock the full features of the Lite version. If you decide you want to upgrade to the Standard Edition, then you can use this same registration code to obtain a £15 discount.

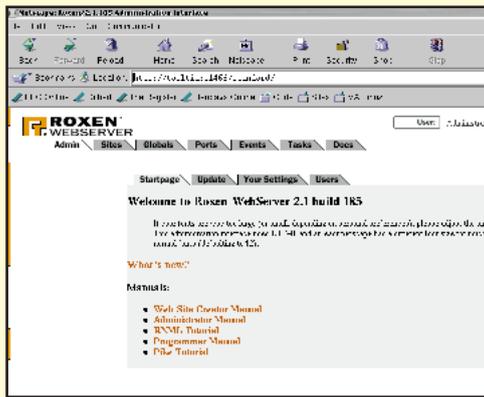
Omnis SYSTEM REQUIREMENTS

Pentium CPU or better 64 MB memory. At least 70MB free on your hard drive. Omnis also requires Netscape 4.x for using the web client and a PDF reader for reading the manuals.

Roxen

The Roxen webserver is the GPLed part of the Roxen Platform - a comprehensive web management system. The webserver itself has been released under the GPL, so you can use it freely. While many people automatically think of Apache when it comes to webservering, Roxen has quite a bit to offer.

For a start it includes its own scripting language and tag system, making otherwise complicated tasks a



Check out our Roxen feature on page 52, and the install it from the CD

lot easier for the web coder. The RXML tags include a special <Gtext> tag which will allow you to render text in a variety of different, cunning, graphical ways. Things like interaction with SQL databases are also made significantly easier.

There is no need to remove Apache just to try out Roxen. Both will cohabit on a single server. For more information on the Roxen webserver, how to set it up and use it, please see the

feature starting on page 52. You'll find the installer in /packages/Web/Roxen on the CD

Mozilla

The world of web browsers is a confusing topsy-turvy place. First there was Mosaic. Next, Netscape and Microsoft's Internet Explorer battled it out to become king of browsers, and in the process opened up a whole world of differing standards and incompatibilities.

The Mozilla project was conceived to change all that. Written from scratch and incorporating the most advanced features, the Mozilla code was built by co-operation between Netscape and the open source community, to create a fast, efficient, feature packed and above all standards compliant browser.

It may have taken a while, but the Mozilla project has finally borne fruit (see our King of Browsers feature in LXF issue 9) in the form of Netscape 6, which was built on the Mozilla code.

Mozilla is by no means finished, as things are being fixed and added all the time. On the CD we do have, however, the latest milestone build of Mozilla - M18 - which is fairly stable (well, more so than Communicator 4.7 on Linux anyway) and includes many of the exciting features you may have read about.

For interested developers, the M18 source code has been provided as a tar file. If you just want to play with it, there are two archives on the CD, one optimised for PPC users, and one for i686 equivalents (ie Intel Pentium II or better).

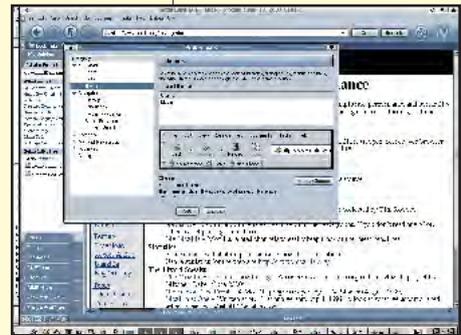
To extract the i686 version, →



Mozilla with the modern chrome looks a lot like Netscape 6.



The Sidebar is an excellent innovation that makes web browsing even easier.



The first thing you'll probably want to do is change the chrome and font settings.

→ open a terminal and try the following:

```
cp <path to
cd>/packages/Web/Mozilla
mozilla-i686-pc-linux-gnu-sea-
M18.tar.gz /usr/share
cd /usr/share
tar xvzf mozilla-i686-pc-linux-
gnu-sea-M18.tar.gz
mozilla-installer/mozilla-installer
```

The *Mozilla* install program will now run. After accepting the licensing agreement, you can choose which components to install and where.

If you have a particularly large (or small) display area, you can change the font settings from the Edit>Preferences menu

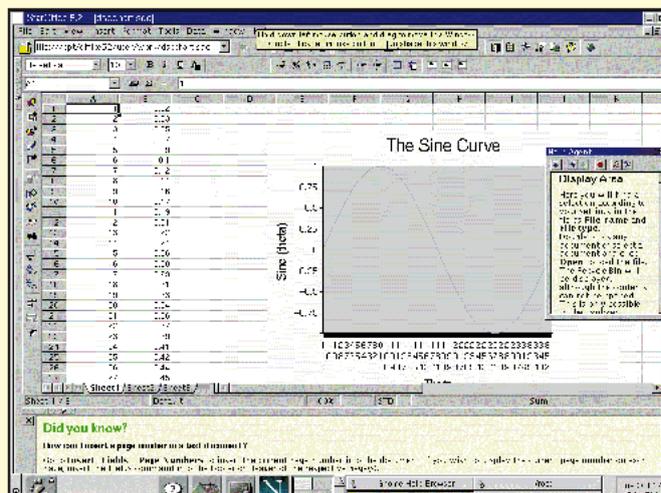
Then it might be a good idea to remove the install folder:

```
rm -Rf mozilla-installer
```

and run *Mozilla*

```
/usr/local/mozilla/mozilla
```

When the application first starts up, you may get the feeling you have seen it



Update source code for OpenOffice appears in the development directory.

before somewhere. The default theme is very much *Netscape 4*, so to avoid unpleasant memories, you may wish to proceed directly to the Edit>Preferences menu item.

The layout of this section is pretty



CUPS is a leap forward in printing, and it's on the CD (as is Omnis).

much the same as Netscape too, so it may be vaguely familiar to you. If you want to try the groovy new “Chrome” (as *Netscape/Mozilla* themes are known), click on the Appearance arrow and choose Theme. Here you can select the ‘Modern’ theme and make *Mozilla* look like, er, *Netscape 6!*

If you have a particularly large or small display area, you may also wish to change the default Font settings (two sections above Themes).

Returning to the main *Mozilla* display, one of the major new features of this browser (and Netscape 6) is the My Sidebar window. This panel resides by default on the left hand side of the screen and is constantly filled with information and content of all sorts. It can be used for searching, to give you a quick list of the links from the page you are currently viewing, and the tabs are completely customisable. To find out more about this feature, the best place to visit is Netscape’s own “My Sidebar” page at: <http://home.netscape.com/sidebar/whatwhy.html>

We hope you enjoy using *Mozilla*, and we will obviously keep you up to date with new releases on the CD, but please remember that this software is not finished! Don’t go calling us saying “this software I got from your CD crashed when I tried to do such and such” We don’t want to hear it. If you think you may have found a genuine bug, surf on over to www.mozilla.org and look up the section on reporting bugs!

CUPS

If you’ve read our What on Earth feature, then you’ll know what CUPS is all about. If not, then read it now! This breakthrough in Linux (and Unix) printing is available to all on our CD.

You’ll find the relevant files in `/packages/Printing/Cups/`, including an RPM, a `.deb`, the tar-ed binaries (for Intel processors) and the full source for everyone else. Once again, sorry if you are a PPC user or whatever, we do try to put as many packages on as possible, but we are restricted by space, and what packages are available in the first place.

The source should be straightforward to build. Simply copy the archive to your `/usr/src/` directory (or wherever seems appropriate) then unpack it and install as follows:

```
tar xvzf cups-1.1.4-source.tar.gz
cd cups-1.1.4
./configure
./make
./make install
```

The basic system includes support for some HP and Epson printers. If you have a different printer, you might like to install the CUPS compatible drivers supplied with the *Gimp-print* package, also in the Printing directory on the CD.

Open Office

We gave you the first release version of *OpenOffice* back in issue 8, but we’ve included this version on the CD as it includes some updates. There are no binary files on the CD this time, just the source code.

Really, this package is only of interest to people who want to keep tabs on how the *OpenOffice* project is progressing, or for developers who might want to check out the code, or maybe contribute to the project. Simply put, if you don’t know how to install and compile this code, it probably isn’t for you! *OpenOffice* is not the most feature packed and stable of Linux Office suites yet, so for everyday desktop use, you may be better off sticking with something like *StarOffice*, which this project was built on.

If you do want to play however, you’ll need to copy this tar archive to somewhere suitable on your hard drive and unpack it with `tar xvf openoffice609.tar`. Then read the enclosed documentation. You’ll find the source code for *OpenOffice* on the CD in `/packages/Developer/Open_office`



Lovely, beautiful scenery. It's just a shame there's a war on. Heavy Gear II is a game of contrasts.

Heavy Gear II SYSTEM REQUIREMENTS

In order to run the heavy gear 2 demo, you'll need some up to date kit, including:

- Pentium running at 233MHz or better
- A graphics card with hardware 3D acceleration, Xfree86 3.3.5 or later, and capable of a 16-colour 640x480 screen
- 64Mb of RAM (though 128 would be better)
- OSS compatible sound card
- \$x CD ROM drive

Heavy Gear II demo

There's nothing like the clash of cold steel against, er, advanced laser weaponry, so prepare yourself for robot mayhem with our *Heavy Gear II* demo. Apparently you'll also have to use your guile and wits to get behind enemy lines, but as far as we were concerned, this demo is about blasting, running and blasting some more!

To install the demo, copy the hg2-demo file from the cd (/packages/Games/Heavy_gear) to a suitable place. Due to a few difficulties, this file may not show up with the correct permissions. To be sure, make sure you are logged on as root and type the following:

```
chmod ugo+rx hg2-demo-
x86.run
./hg2-demo-x86.run
```

The archive will now unpack itself into a directory called hg_2demo. Change to that directory and type:

`./hg2_demo` and the game will launch. You can fiddle around with the options, but the main game modes aren't all available (obviously). There are two 'Historical' scenarios you can play though, both of which should give you a flavour of what the game has to offer, and get you used to the heavy weaponry on offer. Is it, as LXF's battle hardened reviewer claims, the best game on Linux so far? Or do you know better? Send us your opinions at the usual address!

Shogo Demo

For a little variety, we also have a demo of Hyperion's new game (also reviewed this issue), *Shogo*. It's about, er, robots running amok and killing each other. Anyway, lets get to the action. Copy all the files somewhere suitable and execute the self-extracting archive with: `./ShogoDemo.run`

Now, we have some patches for this demo, so extract them with: `tar xvzf ShogoDemoUpd.tgz`

Some new files will be created.

READER REQUESTS

New rules - you ask for it, you get it. That's right, we've started a cunning new service here at *Linux Format* to save you time, effort and money (aren't we nice). Basically, if you have always wanted a particular piece of open source software, but can't afford the download time, we'll download it for you and stick it on the CD! All you have to do is tell us what you want and why, and we'll get it for you. There will be space limitations for this, and obviously you're request will have to be reasonable. Bear in mind that some stuff, although it is freely downloadable, isn't necessarily freely distributable, so we may not legally be able to bung it on the disk for you. Anyway, to make a request, just email us at linuxformat@futurenet.co.uk and put Reader Request in the subject line.



Copy them to the ShogoDemo directory with:

```
cp client ShogoDemo
```

```
cp *.ren ShogoDemo
```

and type 'Y' and return when asked if you want to overwrite files. Now you can run the demo with:

```
ShogoDemo/shogo
```

A popup screen will appear for you to adjust settings. Just click on Launch Game, choose single player and off you go on the City of Ghosts demo level (watch out for the tanks). Happy blasting!

Meet interesting robots and kill them with your massive guns.

Shogo SYSTEM REQUIREMENTS

- Pentium MMX 200MHz with 3D hardware or PentiumII
- 233MHz for software 3D rendering (hardware recommended)
- 32Mb Ram (more recommended)
- OSS/Alsa sound card
- Glibc 2.1/ gtk 1.2

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

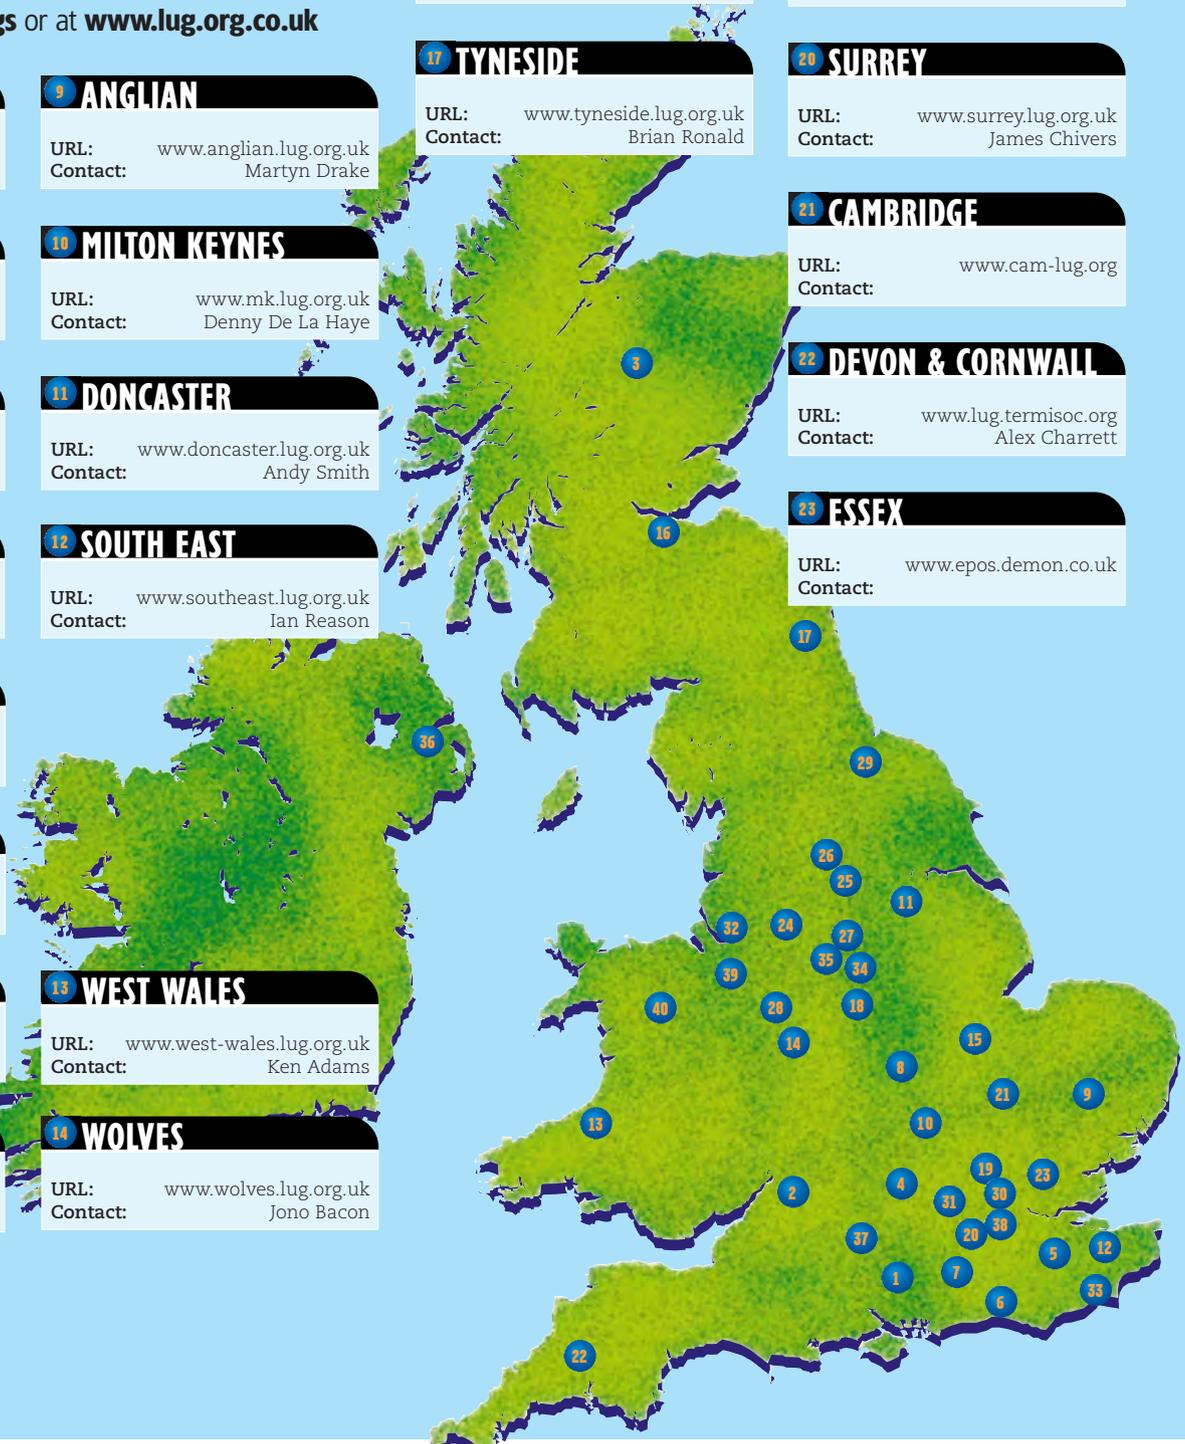
19 GREATER LONDON
URL: glug.linux.co.uk
Contact:

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

25 WEST YORKSHIRE

URL: www.scs.leeds.ac.uk/wylug
Contact:

26 WEST YORKSHIRE

URL: www.wylug.lug.org.uk
Contact: Nigel Metheringham

27 SHEFFIELD

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

28 STAFFORDSHIRE

URL: linux.ukweb.nu
Contact:

29 NORTH-EAST

URL: www.shofar.freemove.co.uk/NELUG

30 LONDON

URL: www.lonix.org.uk
Contact:

31 THAMES VALLEY

URL: www.sclug.org.uk
Contact: Nick Lambert

32 LIVERPOOL OpenSource

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

33 DEAL AMIGA CLUB

Email: superhighwayman@hotmail.com
Contact: John Worthington

34 CHESTERFIELD

Email: info@spirelug.org.uk
Contact: Paul Sims

35 SOUTH DERBYSHIRE

URL: www.sderby.lug.org.uk/
Contact: Dominic Knight

36 BELFAST (BLUG)

URL: www.linux.ie
Contact: Ken Guest

37 WILTSHIRE

URL: www.wiltshire.lug.org.uk
Contact: Jason Rudgard

38 SOUTH LONDON

URL: www.b-Lug.org
Contact: Dr. Tim Traveres

39 CHESHIRE

E-mail: enquiry@sc.lug.org.uk
Contact: Richard Smedley

NEW

40 NORTH WALES

URL: www.northwales.lug.org.uk
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-commitee@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/slugin/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

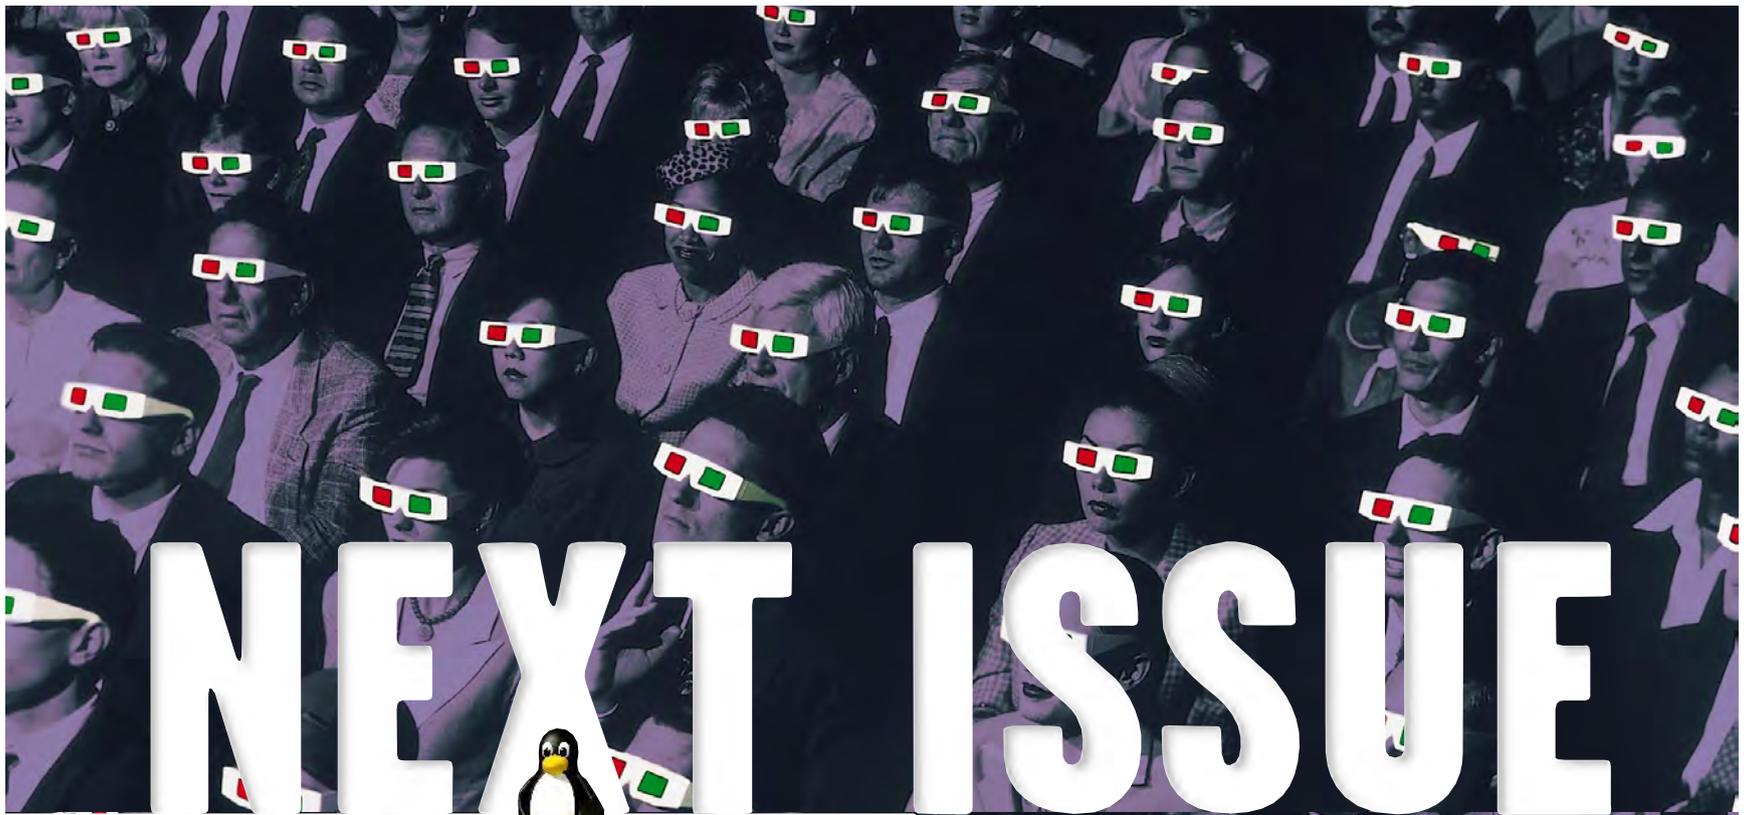
SOUTH AMERICA

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



NEXT ISSUE



ENTER THE THIRD DIMENSION...

Increasing desktop processor power, custom designed super-fast graphics hardware and the increasing demand for 'realistic' games have opened up a whole new dimension for computer users in recent years.

With absolute fortunes being spent on developing 3D sequences for games and films, we'll be taking a look at Linux's place in this new world.

From a practical guide to getting your 3D hardware running properly under Linux, to a look at some of the software available for rendering 3D models under Linux, and a look at how Linux is already making an impact in the world of professional 3D, there's only one place to be!

Reviewed and Rated

Next issue we'll also have an in depth look at SuSE Linux 7 PPC – is it keeping up with its Intel counterpart, and will it keep G4 owners happy? We'll also be having a look at the latest version of Turbolinux, the amazing *Win4Lin* Windows under Linux emulator, and a chance to catch up with Borland Inprise's flagship Java development tool, *JBuilder 4*. Not to mention the latest selection of hot books and open source projects!

Explained

If you want to know even more about Perl, learn how to manage tasks and multithread your Java applications, customise your windowmanager and more. If there's something you'd like to see a tutorial on, just let us know!

Issue 11 On sale Wednesday 31st January

The contents of future issues are subject to change. We're not psychic, man, so don't get upset if some of this stuff doesn't happen!

