Keeping you in the loop

From a recent survey carried out in Germany, a surprising 23% of computer users said they would consider switching to Linux next time they were upgrading their computer systems. And the major reason they cite is the stability of Linux. Perhaps we shouldn't be so surprised, though in the wake of revelations of further security flaws being exposed in Windows2000, perhaps we might have expected security to be a major factor in that decision.

But there is no room for complacency on that score either. Linux is not impossible to crack, nor is it immune to viruses. At the moment we may find comfort in the fact that perpetrators of such anti-social behaviour can find easier targets than a well-maintained Linux box, but the dangers are still there, and if we ignore them and kid ourselves that they don't exist we're only going to make such attacks ever more devastating. Find out what leading expert Graham Cluley has to say about Linux viruses on page 88.

But it's not all doom and despair, we've manage to pack a lot more in to the second of our new-look issues, with a look at how to leverage database technology to your advantage and how to choose which one you should be using on page 42. Tutorials on running multi-headed displays, getting your scanner working and customising *WindowMaker*, along with a review of a game which *isn't* from Loki!

Andrew Arensburger has provided us with a useful talking point in his code commenting feature on page 72, and with reviews of Progeny, Solsoft *NP Lite*, *Photogenics* and more, there's something for everyone in this issue.

Finally, don't forget the CD. We had so many requests for the Sun JDK that we managed to finally wear them down and fill in all the paperwork to let us give it to you, along with the excellent Forté Community Edition – if you want to start using it straight away, check out the Java tutorial on page 78!

Nick Veitch EDITOR

Virus expert Graham Cluley gives us the low-down on the current threats to Linux p88

Don't miss our review of the top-class Photogenics p32

Richard Stallman - madman or messiah? Read our exclusive interview p14

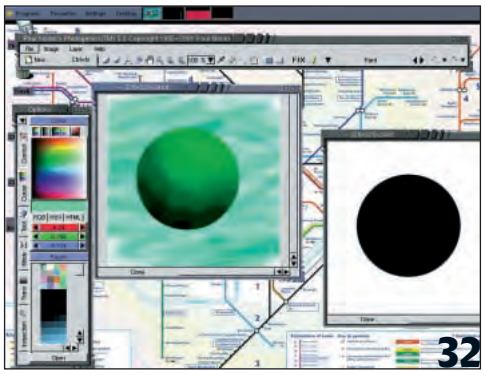

Meet Linux Format's team of writers.....

Andrew Channelle
Our resident newshound, he spends his time alternately chasing news and late contributors.

David Coulson
An *LXF* regular, David is a sysadmin with VA Linux and OSDN. He's also a bit of a Linux security guru.

Richard Drummond
As well as writing our Java series, Rich co-ordinates most of the reviews in the mag.

Paul Ravening
He's the cheerful chap who serves up 650MB of Linux goodness to you every issue. He also has big speakers.

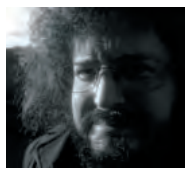

Charlie Stross
Master of Perl, Charlie has been writing about Linux for more years than anyone can remember.

Jon Kent
In the hot seat for Hot Picks this month, Jon has been rounding up the open source finest for you.

John Walker
The only member of the team ever to interview Richard Stallman and survive.

Simon Goodwin
A hardware druid in more ways than one, Simon is currently researching every emulator known.

Jono Bacon
Founder of Linux UK and big *KDE* exponent, Jono is studying Multimedia at Wolverhampton Uni

Biagio Lucini
You'd never guess but he's Italian. An expert on window managers, he is currently studying in th UK.

Aims of the 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:
lxf.subscriptions@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 732398
More contact info on page 97

The UK's best selling Linux magazine

LINUXformat

www.linuxformat.co.uk

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

MASTER LINUX NOW!

COVER FEATURE

20 pages of tutorial, tips and ideas to bring the guru out in you.

Database knowhow

Linux is blessed with two excellent open source databases. But which is right for you? Charlie Stross sifts through the FUD.

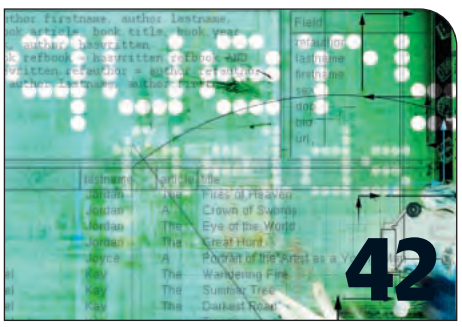

Richard Stallman

During his visit to the UK for the Code Conference, we caught up with the Godfather of GNU for an **EXCLUSIVE** chat about bugs and freedom.

SOAP revealed

SOAP will form the basis of Microsoft's next generation of applications. So, will this cross platform standard mean *Office 2003 for Linux?*

» REGULARS

» TUTORIALS

» REVIEWS

» FEATURES

LINUXformat Coverdisc

Over 640Mb of software to grace your Linux box 90

Give you wallet a break. We separate the wheat from the chaff, bringing you the best Linux applications and utilities without piling the pounds on your 'phone bill

On the CD

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

LINUXformat

In your favourite sections this month...

» DESKTOP

Everything for the desktop Linux user

» DEVELOPER

Hacker tips, tools and techniques

» PROFESSIONAL

Expert advice and info for Enterprise users

LINUXformat Newsdesk

Everybody, it seems, has got something new out this month: We've got shiny new stuff from Red Hat, Mandrake, Loki, Samba, TurboLinux... Could it be spring?

America's most popular distro gets a new lease of life

Mandrake 8 out now!

Mandrakesoft are the latest vendor to release a kernel 2.4 based distro. But what other goodies have been cooked up for version 8?

After a short gestation in the famed Cooker, Mandrake 8.0 hit the mirrors on April 19th, promising improvements that would benefit beginners and experienced users alike. The company said this release confirmed its 'commitment to offering the most complete, powerful and user-friendly' Linux operating system.

As well as the standard download version – available from www.mandrakesoft.com – Mandrake 8.0 will be available in three different retail packs: Standard – for desktop and basic server use. Powerpack – for workstation and server use. ProSuite – for enterprise users.

Frederic Bastok, Co-Founder of MandrakeSoft, said the release was part of the company's strategy to 'release soon, release early' in the spirit of Eric Raymond's The Cathedral and The Bazaar. "Our approach of making the distribution available on the Web before the launch of the commercial packages is an effort to thank the numerous developers and testers who constantly contribute to this fabulous project," he said.

In an effort to promote ease of use, Mandrake 8.0 includes the latest releases of both main desktop environments (*KDE 2.2.1* and *GNOME1.4*) and also features a new, improved DrakX Installer, which builds on its enviable reputation for hardware detection. Users will also discover that maintaining the system is easier thanks to the all-new control centre which gives a very Windows-like view of your hardware.

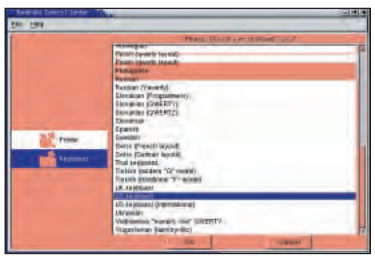

Mandrake 8 promises to make configuration a breeze.

"For the first time, Mandrake will default to a 2.4 series kernel"

***Evolution* is a new calendar app.**

For the first time, Mandrake will default to a 2.4 series kernel to take advantage of the wider USB support, better SMP (Symmetric MulitProcessing), brand new IP and firewalling stacks, and 64Gb memory support. However, users will be able to 'step back' to a 2.2 series kernel if compatibility issues arise.

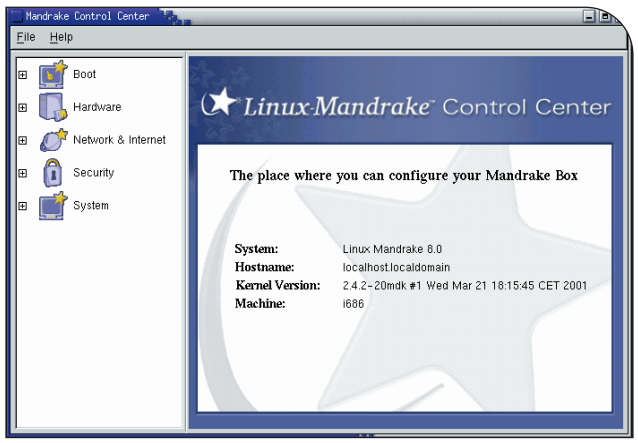

The control centre is very Windows-like.

What's in the box

As ever, the main part of the distribution is augmented with a wide range of extra applications. "It comprises over 2,300 applications, including *StarOffice* and IBM *Via Voice*, the market-leading voice recognition system" Bastok said. While all four versions feature a range of general administration and productivity software, the ProSuite edition incorporates a selection of server tools with regular security updates included in the initial purchase price. These tools include an optimised version of *Apache Web Application Server*, to deploy Web applications, POP3 and IMAP for reliable email servers and a selection of the most efficient database servers. It will also feature graphical configuration tools to make the most of the latest kernel's security and firewall features.

For the more leisure-inclined, the distros will also come with a massive selection of games including the excellent open source flight simulator *FlightGear*. And, to make sure you have enough time to play all these games, there's a preview release of the new *GNOME* calendar/scheduler application, *Evolution*.

Early reaction among reviewers and proper users suggest that Mandrake 8.0 is a real improvement over previous versions, but if you want the definitive view, look out for our assessment in issue 16.

Concept handheld

Project Mercury

For months we've been tantalised with the idea that you can, with a little bit of hacking, get Linux running on an iPaq. In fact, the hack even has official blessing from Compaq (see www.handhelds.org for more), but if you wanted an iPaq without Windows off-the-shelf, well, you were out of luck. This may all change in the near future. www.infosync.no, a Norwegian news website, claims to have uncovered Compaq's Mercury project.

Project Mercury is a back pack (or BackPAQ) for the colour H3600, which will form the basis of an MIT initiative on pervasive computing, Oxygen.

The rumoured feature-list includes the usual PDA functions, a pair of PCMCIA card slots, an audio interface with headphone jack and an accelerometer which, it's suggested, could be used as part of the screen navigation system

Defining Linux

Creating the standard

One of the barriers stopping widespread adoption of Linux has been the perceived differences between the various distributions. Although the competitors *are* Linux, there has been no standard way of installing software or organising file systems. As well as being confusing for users, this has also meant that software developers and publishers need to provide installers and binaries for a large number of competing distros.

In order to combat the problem, advocacy group Linux International set up the Linux Standards Base group with the intention of defining a uniform platform that developers and distro vendors could work towards. This standard, it was hoped, would cure some of the problems that users face when choosing and using Linux software.

Through this long gestation – and the demise an ill-fated competitor (the Linux Standards Association) – the group have been working to define the standard and set up adequate compliance testing. One of the major landmarks in the development of the standard has been the recently released Linux Development Platform Specification, which was created to allow developers to create portable code that will work with all compliant distributions. The document's introduction states that the specification is designed so that programs developed on a conforming platform are expected to be portable to all generally available Linux distributions as of March, 2001.

Daniel Quinlan, editor of the LDPS said release 1.1 of the specification demonstrated a real commitment to help both Open Source developers and companies to develop portable applications for Linux. "Open Source development moves very quickly, so it's important that the LDPS be kept up-to-date until the more stabilising Linux Standard Base is released."

Most of the popular distros currently meet this specification, but there are still major differences in the locations that the various systems store their files. This problem should be address by version 2.2 of the Filesystem Hierarchy Standard (FHS) which should allow software to predict where installed files and directories by specifying 'guiding principles for each area of the filesystem'.

The FHS is currently implemented in a number of distros including the most recent releases from Caldera and SuSE. In fact, Caldera have become the first mainstream vendor to offer a distro that conforms to all aspect of the LSB, using a service pack which installs over the latest beta version of OpenLinux.

http://www.freestandards.org/ldps/
http://www.linuxbase.org/

Major changes

Linux Development Platform Specification 1.1 has incorporated a large number of changes:

- **Printing documents in portable applications.**
- **New POSIX threads information.**
- **An expanded "Frequently Asked Questions" section.**
- **Replacement ncurses specification.**
- **Additional information on RPM.**
- **Notes on securing applications.**

NewsBYTES

SuSE have announced a new program which will see SuSE 7.2 installed in 2000 public and private schools in the USA. Dirk Hohndel, president of the company's stateside operations, said the initiative was aimed at showing students, administrators and

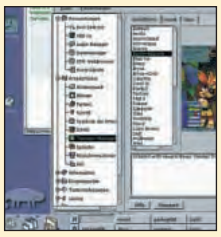

teachers that they don't "have to be tied to expensive operating systems and costly upgrades."

WEBppliance LS is a new hosting appliance software product from Ensim Corporation. Designed for servers running Red HAT, WEBppliance features three administration levels to provide management facilities for service providers, resellers and end users. www.ensim.com.

Embedded computing firm **Wind River** have purchased Berkeley Software Design Inc (BSDi) the commercial home of FreeBSD and Slackware Linux. Announcing the acquisition, the company said: "Wind River will continue to support the efforts of FreeBSD open source software, working closely with the open source community of developers." The purchase didn't involve Slackware which will remain with the newly named iXsystems.

Executives from the world's biggest computer companies will be getting together in October to discuss ways to bridge the digital divide. But, in what the organiser have said was an 'oversight,' no one from the Linux community was invited to the event.

Obviously the organisers hadn't seen the latest projections from IDG which suggests that Linux is rapidly gaining popularity in Latin America, and will be running on up to 33 per cent of the region's desktops by 2003. Linux is also expected to become the dominant operating system in China and India.

Easier networking

Samba 2.2.0

The Samba Team have announced the availability of version 2.2 of their award-winning implementation of the Microsoft CIFS/SMB protocols for Unix systems. The software allows the easy integration of Linux/Unix and Windows systems into one unified network.

The latest version is billed as a significant update – brought about in part by Microsoft changing (or breaking) its own network protocols in the latest Windows releases – and has been tweaked to take advantage of some of the features of the latest kernel.

Improvements include:

- Full implementation of the Windows NT point-and-print functionality independent of Microsoft code.
- Integration between Windows oplocks and NFS file opens.
- Ability to act as an authentication source for Windows 2000 and NT4.x clients.
- Integration with the winbind daemon to provide a single sign on facility for Unix servers in Windows networks.
- Support for native Windows 2000/NT4 printing RPCs.
- Support for server supported Access Control Lists (ACLs).

Link: http://www.samba.org

Suite update

KOffice beta

The beta process for *KOffice 1.1*, the integrated office suite for *KDE*, has begun in earnest with a projected stable release for the summer. The developers' claim the primary goal of this release is to allow users to test and request new features, and code development is concentrating on stabilising the various applications and ensuring they operate in a unified fashion.

Big changes are planned for the flagship application, *KWord*, which will gain more Desktop Publishing style tools. However, the biggest changes should be in the program's ability to deal with very large documents, improved find/replace tool, better table support and a smaller file size.

Other applications to see a major upgrade in this release include *Kpresents* (which gets the ability to embed other parts of the suite in documents), *Kspread*, *Killustrator* and *Kivio*.

Krayon, which was formerly known as *KImageShop*, will be included in the office suite for the first time and will offer professional level photo manipulation. Other new additions include *Kivio* (flowcharting), *Kchart* (charting) and *Kugar* (reporting).

New releases and a great offer

Out now! Tribes 2 and Alpha Centauri

Loki games have finally shipped two to the most eagerly anticipated Linux games, and offered gamers (and LUGs – of which more later) the chance to get both games at a knock down price.

Tribes 2 is a gloriously mad first person shoot 'em up in which players don giant armour suits before joining (or creating) a 'tribe' with which to beat up all the other teams and take control of worlds. Tribes can consist of up to 60 human or AI controlled members and, unlike most FPSs, the key to coming out on top in *Tribes 2* is cooperation.

For multi-player gaming over a network – including the Internet – *Tribes 2* offers a number of ways to stay in contact with your team-mates, including email, voice chat, instant messaging, message forums and websites. There are also a number of training missions to ease you into the concept of online gaming before even warming up your modem.

Tribes 2's 3D engine is simply superb, and the gameplay's okay too.

Tribes 2 uses a completely new graphics engine and includes an innovate terrain editor that lets you create new levels 'on the fly' which should increase the longevity of the game beyond the usual few months.

Beyond Civ

Alpha Centauri is the classic strategy game from the creator of *Civilization*, Sid Meier. *Civilization* saw you take on the mantle of a great leader attempting to nurture your followers from small-town agriculture to world domination, all the while smiting your opponents wherever possible. Alpha Centauri takes this concept to its logical conclusion, with players racing to become the first to leave the overcrowded planet Earth and conquer new worlds.

The Loki port of *Alpha Centauri* includes not just the original game, but also the Alien Crossfire add-on pack.

Both games were released in mid-April and should be available from your favourite online seller. However, if you buy online from Loki, you can pick up both games for under $60.

Interiors locations are very atmospheric, but things can get out of hand in confined spaces.

Tribe 2's exterior shots look pretty cool as well.

While not as graphically stunning as Tribes 2, Alpha Centauri boasts incredibly addictive gameplay.

LUG offers

Loki have announced a new scheme which will allow Linux User Groups (LUGs) to buy their most popular games at discounts of up to 50% – American LUGs don't even have to pay the postage!

Loki's Kayt Sorhaindo told us that User Groups from around the world could participate in the scheme, but users outside the US would have to pay shipping charges. Details of how to order can be found on the Loki website:

www.lokigames.com

The games available are:

- **Sid Meier's Alpha Centauri Planetary Pack and Tribes 2 Bundle – $29.95**
- **Tribes 2 – $24.95**
- **Sid Meier's Alpha Centauri Planetary Pack – $14.95**
- **Eric's Ultimate Solitaire – $9.95**
- **Heavy Gear II – $14.95**
- **Heretic II – $14.95**
- **Heroes of Might and Magic III – $14.95**
- **Myth II – $14.95**
- **Quake III Arena – $24.95**
- **Railroad Tycoon II Gold – $14.95**
- **Simcity 3000 – $24.95**
- **Soldier Of Fortune – $24.95**
- **Descent3: – $12.45**

SELinux latest

PGP/NSA join up

America's National Security Agency (NSA) have announced a £1.2million dollar deal with NAI Labs, a division of PGP Security, to further develop Security-Enhanced Linux, a prototype ultra secure Linux distro.

NIA's work will focus on reducing the risk of security breaches caused by flawed or malicious applications. Both organisations have a history of developing controls and protocols for improving security on open source projects.

Terry Benzel, NAI Lab's Vice President of Security R&D said the work would help to create and protect a secure Internet infrastructure. He said: "Open source platforms provide the basis for developing online communication and business strategies, and our work ensures that the building blocks are secure." The two groups, Benzel said, had identified a need for robust security policies in mainstream operating systems, yet most "lack critical security features."

NSA initially chose to develop their secure OS with a Linux base due to its growing adoption and because the open development environment provided an opportunity to show that a mainstream OS could also be secure.

http://www.nsa.gov/selinux/

Mysterious and spooky

The National Security Agency is America's own cryptology agency and their mission is to "understand the secret communications of our foreign adversaries while protecting our own communications." The organisation is thought the biggest employer of mathematicians in the US and runs the The National Cryptologic School, which is training the next generation of code-breakers and makers.

The NSA began the SE-Linux project to "meet the security needs of a wide range of computing environments."

NewsBYTES

The **Simputer** (which stands for Simple, Inexpensive and Multilingual) is a low cost Linux-based computer designed to bridge the digital divide in India. Retailing at about the same price as a colour TV ($200), the Simputer's makers hope the device will improve opportunities – and Internet access – for India's 995 million non-computer owning citizens.

Just three months after releasing the source code, developers of **OpenDivX** – the video compression technology which allows delivery of full screen, high quality video over the Internet – claim the system has been ported to all the major operating systems including Linux, MacOS and Windows. The core technology is also available for a number of embedded processors. The project has 1,500 registered users and over 100 active developers.

Bill Gates has slipped one place in the annual Sunday Time Rich List. Robson Walton, boss of giant retailer Wal-Mart is said to be worth $45.3 billion, while Gates has to make do with a paltry $37.5bn. Other geeks with more money than can be imagined were Larry Ellison (3rd), Paul Allen (6th), Steve Ballmer (15) and Michael Dell (25th).

Odd port of the month has to be **Sony's** decision to bring an official version of Linux to the Playstation 2. The distro is based on Red Hat and uses the console's PC Card slot to attatch the 40Gb hard disk, keyboard and mouse. The company plans to ship 1000 of the units to test the water. The bad news is: it won't work with European systems.

US developer **Bynari** have announced that their new email/messaging client will interoperate with Outlook without any form of proxy. This, the company claim, will "open up a world of new opportunities for integrating the Linux and UNIX desktop user community into the Windows-centric enterprise."

More new releases

The new distros just keep on coming!

TurboLinux and Red Hat have both joined the distro rush, releasing Server 6.5 and version 7.1 respectively.

TurboLinux Server 6.5 is aimed squarely at the enterprise market and features extensive support for five languages. It also introduces an innovative mass installation tool.

Craig Oda, Vice President of Product Management at TurboLinux said there was a real need for multilingual support. "This software demonstrates our focus on simplifying Linux adoption across multiple languages," he said.

A trail version can be downloaded from **www.turbolinux.com/downloads/.**

Red Hat's new distro is the latest to be built around the 2.4 kernel, and features a number of improvements aimed at enterprise users. Red Hat 7.1 also has a collection of new configuration tools which, the company claim, should make setting up and administering DNS, web or print servers easier and tight integration of the Red Hat Network system that should simplify the process of updating software.

The release, said CEO Matthew Szulik, would demonstrate the obvious benefits of open source software to organisations and business users. He said: "The elements in Red Hat Linux 7.1 work together to deliver the most powerful, automated open source operating system for fast-growing enterprise and Internet infrastructure users."

Red Hat have release version 7.1 of their distro.

IP over avian carrier

Penguins and pigeons

Some people just have too much time on their hands. 11 years after its proposal by David Waitzman (on April 1, 1990), the Bergen Linux User Group (BLUG), have finally found a way of implementing the Network Working Groups IP over avian carriers (rfc1149) protocol.

The initial proposal stated that 'Avian carriers can provide high delay, low throughput, and low altitude service.' with the advantage that 'unlike some network technologies, such as packet radio, communication is not limited to line-of-sight distance.'

Using the methods set out in the protocol (see **http://www.rfc-editor.org/rfc/rfc1149.txt**), BLUG marshalled the services of a number of pigeons to begin their experiment on April 28 under the watchful gaze of Linux kernel guru Alan Cox.

'We decided to do a 7 1/2 minute interval between the ping packets,' the report said, 'that would leave a couple of packets unanswered, given ideal situations.'

A BLUG member prepares the delicious, plump breasted messenger.

However, the conditions turned out to be less than ideal and the first data packets were rerouted after neighbouring pigeons turned up to play. The report continues: 'after about an hour of fun, we could see a couple of pigeons breaking out of the flock and heading in the right direction. There was much cheering.' The carriers eventually arrived at their destination.

The ping took 6165731.1 ms.

Now online at the Linux Format website.....

Ease of use

One of our correspondents this month has suggested that making Linux easy to use will involve loosing its functionality. Is it the case that comprehensive means complex or can you combine simplicity with power? Is Linux being dumbed down to appease the masses? www.linuxformat.co.uk

HP's missing link

Hewlett-Packard are putting their weight behind an application server that they claim will provide the missing link – alongside Apache and Linux – in the open source server architecture.

The Enhydra 3.5 application server has a low profile past, but its adoption by HP with their massive worldwide sales force virtually ensures a high profile future.

The deal will see HP resellers and distributors recommended Enhydra as an approved application for HP platforms.

Apache, Roxen and...

The third way?

The nice, simple open source database world – see feature on page 42 – could be turned on its head by the arrival of the newly opened *SAP DB*, an SQL-based relational database management system. *SAP DB* has been in active, proprietary development since 1994, but in October 2000, its developers used the LinuxWorld Expo to announce a timetable to open source the technology, with the database engine being released under the GNU General Public License (GPL), while the Lesser GPL license would be used for the system's clients and application programming interfaces.

SAP DB, the developers claim, not only provides high availability and performance scaling, but also offers object orientation and the ability to deal with 'unstructured' data. It supports SQL (including Oracle 7 compatibility), JDBC and ODBC, access from Perl and Python and HTTP-based services with either HTML or XML content.

Explaining the change to an open source rather than proprietary model, SAP said "times have changed, that databases are becoming part of the basic technology infrastructure, and as such, they need not be proprietary or complex." The company said they were glad to be part of the open source revolution and that the move would benefit the entire database development community.

The Berlin-based company have set up a site for developers and users with online and downloadable documentation, sample applications and the latest news on the project.

http://www.sap.com/solutions/technology/sapdb/framesets/sap_db_software.htm

Embedded Linux News

» MontaVista have announced the availability of the next version of their professional embedded OS. The company claim the 2.4-based Hard Hat Linux 2.0 combines unrivalled support for embedded processor architectures and a feature-rich suite of embedded development applications. **www.mvista.com.**

» The Embedded Linux Consortium (ELC) has moved to prevent the fragmentation of the industry by proposing a single unified specification for the OS. Board members also floated the idea of an 'Embedded Linux' mark to be used in the same vein as the Intel Inside device.

» Sun's Java developers have been burning the midnight oil trying to adapt the language to make inroads into the one area it doesn't feature: real time operation. Greg Bollela, leader of the team, told delegates at the Embedded Systems Conference: "We're operating at our limits of our ability to comprehend just what we're doing!"

» Korean PDA maker PalmPalm have linked up with Conversay to create the first Linux-based G3 communicator, which will also have extensive voice-recognition abilities.

» With Agenda failing to hit its deadline for availability of the VR-3 PDA, Korea's G-Mate have announced a price cut for the developer edition of the Yopi PDA.

LinuxWebWatch

Report bugs, speed up Mozilla!

KDE's bug wizard is very useful.

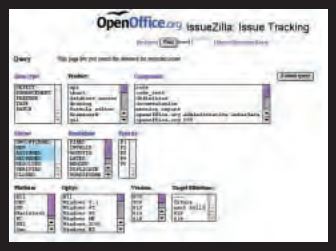

OpenOffice also simplifies the job.

Loki want testers to play games.

Bugwatch – doing your bit

Richard Stallman wants you to put your experiences to good use.

If you've read the RMS interview starting on page 14, you'll know about the man's passion for users reporting bugs. One of the benefits of community development is that hundreds – or even thousands – of users (who may not have the skills or time to actively get involved in a project) can make their mark by testing applications in real world situations, and reporting back problems encountered to those who can make changes. The biggest projects, of course, need the most help.

Mozilla (www.mozilla.org) is the monumental browser project that has formed the basis for a number of other applications including *Netscape 6*, *Jabberzilla* and *Galeon*. The development team make nightly builds and, as well as having a comprehensive bug reporting system (www.mozilla.org/bugs/), the project also has facilities for tracking bugs that you've reported.

An even bigger project – and one which could have even bigger benefits for the Linux world – is *KDE* (www.kde.org). Realising the importance of bug hunting, the team have created a wizard system to make submitting a bug report as easy as possible. You simply go through each step, selecting which application, version etc that you have, then finally input the nature of the bug.

When Sun decided to open source *StarOffice*, creating *OpenOffice* (www.openoffice.org), they chose to use a modified version of *Mozilla's* Bugzilla reporting system. Current areas of concentration include Unicode support, XML file formats and overhauling the user interface (UI). Developers testing an application like *OpenOffice* may throw some light on a few problems, but only real world use will reveal the bugs that real world users will find most annoying.

And it's not all work. Loki (www.lokigames.com) are currently developing a Linux version of the stonking FPS Deus Ex. In order to get it right, they need beta testers who can play through the game and report back any problems they find. The benefits are that you get to play the game before anyone else and, when it finally reaches the shelves, you will have played a part in advancing the cause of Linux gaming.

Loki do tend to bombard testers with emails, but this appears to be a symptom of their desire to really get a game right before it goes 'gold'.

CODE Clash

John Walker dons a hard hat and enters the fray at the First CODE Conference, where free software and open source advocates come together. Let's get ready to rumble!

It was inevitable. Inviting the leading members of the Free Software movement and the leading members of the Open Source movement to the same conference. There was bound to be trouble. Throw in a collection of academics, artists, business men, and thrill-seeking journalists, and you've got yourself a guaranteed Argument-In-A-Can.

And so it was at the tentatively titled 'First CODE Conference', taking place at Queens College Cambridge this April. Major names from an eclectic selection of arts, science and tech backgrounds spent three days telling each other what they thought should be the future of copyright, patent law and all manner of collaborative works.

Despite the array of near-deadly differences, the overriding tones of the week were a universal terror over the concept of software patenting, and an awareness that working with your neighbours may not be so abhorrent after all. Shock, horror!

Love Thy Neighbour

As a reader of a fine journal such as this, you already have a head start on the vast majority of the world. The chances are that if you aren't already running a non-proprietary operating system, you are seriously considering the idea, and hence will already have a brain orientated in a forward thinking direction. The same bias was almost entirely the case amongst the 250 or so attendees of CODE – so rather than preaching to the converted, a lot of time was spent pondering and suggesting how such radical ideas could be presented to the masses. How do you manipulate Mr Big Business to consider that there may be a future in 'Giving Stuff Away'. If ever an idea was to strike daggers of fear into the atheroma of the corporate world, it is this.

The most eloquent testimony was given by entrepreneur Bob Young. Co-founder of Red Hat, and very rich person, he is a man who realised the market potential of free software and got it right. Red Hat is now the most popular distribution of a GNU/Linux operating system, and has a tidy turn-over. As Young would say himself, this was all achieved by "giving our product away for free." He stands as living proof that the ideas endorsed by such a community are not as ridiculous as certain major software companies would have you believe. Does this sound obvious to you? Try telling the person sitting next to you.

But his voice was not alone. Speaker after speaker stood up to tell how embracing an open source business method had advanced their work, saved their company, or even opened the possibility of saving lives. Tim Hubbard, head of Human Sequence Analysis at the Sanger Center told of the race to complete the human genome. It was one of the most significant battles between the public and private world, as open source took on private business in a race to see who would win the rights to the genetic make up of our own species. Drop to your knees and thank your God that Hubbard and his colleagues won in the end, or perhaps you would now be asked to pay for a licence every time you wanted to have an adrenaline rush. While the scale of this situation was very much underplayed, it was becoming clear quite how important such ideologies are becoming today.

Other speakers ranged from Roger Malina, an astronomer from the CNRS-Laboratoire d'Astrophysique de Marseille, and formerly of NASA, who promotes releasing data into the public domain for analysis, to Christopher Keity, an anthropologist studying the effects of free software on a cultural, economic and political level. The key word was diverse.

Turning the Tables

However, despite the majority of attendees being of a similar frame of mind, this did not prevent the ever-present in-fighting from surfacing. As anyone who has communicated with him will know, you don't invite Richard Stallman somewhere and not expect to have a little bit of a show.

His lecture on the Thursday morning was perhaps a reasonably predictable affair. Beginning with a walk through the history of copyright – from its beginnings to its current state – he entered into a speech he has given quite a few times before. Nevertheless, it never does any harm to hear it again, and St IGNUtious went on to discuss the inherent ethics and principles that underline all his actions. It is quite a remarkable experience to listen to a man who is so driven, and so dedicated to his morality; a man who will not compromise at any level, as we were later to discover.

During his hour long talk, he dropped in a few nuggets of controversy, suggesting that mad cow disease, foot and mouth, and the state of Britain's trains could all have been prevented if we weren't a

From left to right: Anthropologists Christopher Keilty and Marilyn Strathern, conference Chair John Howkins, and confused Swede Geert Lovink.

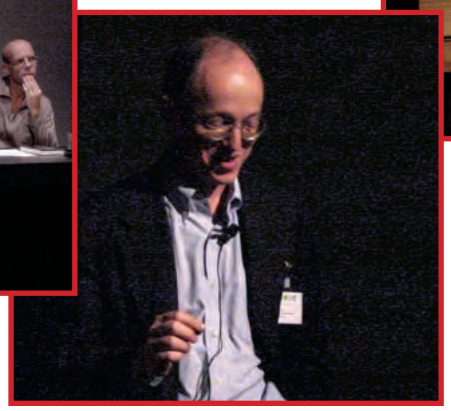

Bob Young talked about "The Future of Innovation", but mostly about Red Hat.

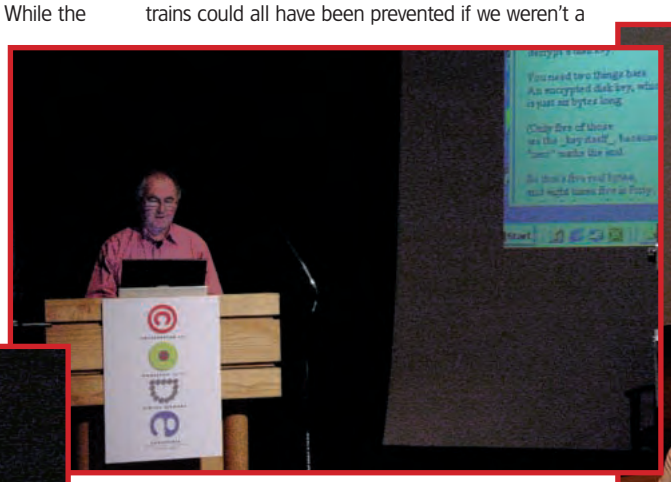

Cambridge Professor of Law, Bill Cornish, broke all sorts of laws, demonstrating various methods of perpetuating the DeCSS code.

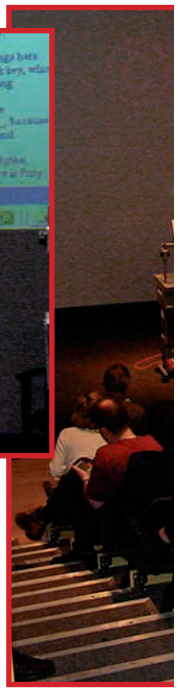

country that insists on refusing to look at the long term consequences of capitalist business. One rather got the impression he believes that were he in charge, none of these things would have happened. Time was given to chastising Stephen King for his abuse of the 'Micropayment' scheme in his recent online experiment, and a request for a complete boycotting of e-books. There was of course the oft-repeated reminder of why we good-thinking citizens should not be saying "open source" or "Linux", delivered with all the passion as if it were the first time he had preached it.

His obscure delivery, coughing and looking lost for the most part, somehow did not irritate as it would were someone else bumbling around in such a way. He certainly possesses a presence that excuses him from being treated like any other human being. In fact Glyn Moody, in his after-dinner speech, compared Stallman to Beethoven, claiming the pair to possess a similar drive and genius. Whether this is true or not, it is certainly the case that they both contain a similar temper.

He reappeared and announced in the voice of a six-year-old girl "I'M LEAVING NOW!"

But what of the all-exciting brouhaha? Late Thursday afternoon, having just been treated to a two hour interview with the great man, your intrepid *Linux Format* reporter made his way back to the lecture theatre to discover Stallman had beaten him to it, and was already causing trouble. Standing next to him in the galleries, it was clear that the lawyer on stage was beginning to rile him. Stallman had unfortunately missed the first half of the talk where the negative aspects of software patenting had been described, and had arrived in time to hear only the suggested positives. Having already been interrupted once, lawyer Justin Watts was perhaps deliberately provoking his latest audience member by repeated uses of the term "open source" and refusing to say "GNU". The final straw was when he cheekily suggested that the Microsoft of the future may have a right to sue the 'open source community' for possessing a monopoly. That was more than Stallman could bare, and bursting forth he grabbed a microphone from an unsuspecting student, and let rip with a screaming fit of abuse and outrage. Eventually shouted down by the event organisers, and having had his mic switched off, Stallman turned and stormed from the room. A reasonably impressive show, that was sadly let down by a rather unfortunate encore. He immediately reappeared and the in the voice of a six year-old-girl throwing a tantrum announced: "I'M LEAVING, NOW!" Long-time friend Bruce Perrins jumped from his seat on the stage to follow him out, following what was clearly a well-worn path.

Perrin's counselling talked Stallman down from claims that 'everybody hates me' and 'they're all against me' (during which Perrins used the phrase, "Everybody is a potential friend Richard"), to the point where Watts and RMS were found to be outside talking it all through rationally. Eye-witness reports found that Watts was coming around to Stallman's way of thinking! These incredible scenes are par for the course when Stallman is around, but still the conversion of Watts was a remarkable result, and testament to the enormous evangelical power of the founder of all things shiny and GNU.

Evening Service

When conferencing, evenings are always so difficult to fill. Your guests have sat in a lecture theatre all day, and will want some sort of alternative stimulus. And this stimulus is if course alcohol. But in between finding tables of free booze, various activities were laid on for the eager participants.

The first night saw one of the most eclectic mixes of humans trudging through the windy Cambridge roads towards the narrowest building ever built. Three flights of quite possibly the most dangerous stairs in the world led to a small white room displaying a cartoon on a wall. Only upon listening to the CGI female on the blue background was the reason for this being revealed. As part of an experiment in weaving the ethos of Free Software into Art, a group had bought the rights to a Manga character, and then open sourced her, so that she was freely available to the world. The film consisted of her having a very calm existential crisis, realising her own lack of reality, and then seeming to be fine with that.

The second night saw the world's first screening of the soon-to-be-released film *The Code*, from Finnish director Hannu Puttonen. It attempted to tell the tale of the evolution of the operating system we know and love, while dismissing some of the myths that haunt us all. Sadly, his decision to ironically present Linus Torvalds as some kind of saint backfired, doing more to perpetuate the idea that he is solely responsible for all the Good In The World, than burst it. Head to head interviews were interspersed with documentary footage of various Linux conferences, and an incredibly Hello-esque look though the Torvalds residence. Fortunately Stallman left halfway through, as a repeat performance did look on the cards.

So what was achieved? In a very poetic way, the sense of community was probably the most important result. It became clear that ideologies that for so long have been restricted to a small group of self-confessed geeks are spreading wildly across all manner of contexts and lifestyles. To realise that the same ethical approach held by the bedroom hacker is also embraced by the world class scientist is to realise that this is certainly something that isn't only going to change the world, but something that already has. And is. Like some sort of -correct Mason-esque society, the open source/free software community is forming a tight and secure network around the world; a safety net perhaps for when the proprietary bubble bursts.

RMS stood Jesus-like on the stage afront all his worshippers.

Another collection of poeple explaining how open source software had advanced their particular field.

Richard Stallman

Hacker, zealot, madman, genius...

At the recent CODE conference in Cambridge, John Walker sat down for a cosy chat with the godfather of free software and all-round hacker guru, Richard M. Stallman.

None of us would be here if it were not for one man. There would be no GNU/Linux, there would be no cute and cuddly penguins, and heaven forbid, there would be no *Linux Format*. And that man is *not* Linus Torvalds.

Richard Stallman decided he was going to change the world. Many of us may have had similar naïve ambitions at one point in our lives, but very few of us can claim to have achieved it. But then very few of us are geniuses.

Genius is a very bandied about word, grossly misused and applied in the most oxymoronic ways, such as to footballers. Here there is no such hyperbole. As a hacker he is without peer. This is a man with an incredible stamina for his passion, and a dedication to a cause that is almost religious. Glyn Moody, writer of *Rebel Code*, compares him to Beethoven, due to his incredible natural ability and conviction to his work. Genius, of course, is never this simple. It is a borderline personality disorder brought about by extreme intelligence, and Stallman's personality certainly has some, well, interesting quirks.

It would not be possible to portray RMS in this interview without recounting an honest description of his actions, but it is very important to realise that his erratic behaviour is borne out his absolute dedication to living the life he preaches. It is also very tempting to make big play of the fact that mid-discussion he will get up and begin his strangely formulaic dance moves, or any of the other famous overgrown nervous ticks. But one begins to wonder whether these are part of the performance. Does he do the tricks because on some level he is expected to, or are they genuine habit?

However, all such details should not be overblown. Richard Stallman is a man with a message, and it is his message that is the star of the show. Such habits are an irrelevance to him, and hence should also be so to us.

The interview took place in the Great Hall of Queens College Cambridge, sitting around an oversized wooden table. There were a few of us there, and thanks to the extreme kindness of *LXF* – allowing a couple of people from student newspapers and the like to join us – the interview became 'open source'. Perhaps centuries of great discussions had taken place here, but I think it is safe to say that on none of these occasions would anyone have been wearing a jumper as awful as Richard Stallman's. How he got it through customs is a mystery. But perhaps it possessed some form of hypnotic property, as within a few seconds we were discussing his life-ethos.

LXF You have an extraordinary passion for your work, and one that seems to override everything else in your life. When did you first realise the strength of these passions?

RMS I came to these conclusions about the time I was starting the GNU project, and that's why I started it.

LXF In your writings, the beginning of the project and the birth of your passions seem to coincide. One must have happened before the other.

RMS Well, I had the experience first of living in a free software way of life, and I liked it. I always had a tendency to think about it in more ethical and political terms than the other people. They also liked the way of life, which is why they continued to live it. But I started to think of this as a matter of principle, and thought: "This is the good way to do things. I want to uphold this." I'm willing to fight or make sacrifices. It's something good.

And so, when that community was wiped out by a combination of various things, I was suddenly faced

with the prospect of participating in the life of using proprietary software. And it looked so disgusting, shameful, ugly that I decided I was going to build another way of doing things.

LXF Did you have that kind of wonderful arrogance from the very beginning? Were you absolutely sure that you were capable of this?
RMS I wasn't absolutely sure I was capable of it, but I was sure I had a chance, and that it was worth making the effort. Was I sure of success? No. But when you are fighting for something like freedom, it shouldn't be a question. If there's a chance that you are going to fail, is that a reason not to try?

LXF But this must have appeared as such a monolithic task in front of you? It must have occurred to you that 'wow, I'm 15 or 20 years away from achieving my goals!'
RMS Well, I was thinking that it might be five years. I didn't know how many people would help me. I didn't need to know the answers to these questions. I knew that I had found something that was worth trying to do. The other things I could have done in software were not worth trying to do. It would not have been good if I had accomplished them. This, I knew if I accomplished it, would be a good thing. It was worth trying to do because it would give people a chance to live in freedom.

LXF You joined the world of computers through this free system. It was free, and then this freedom was taken away from you?
RMS Indeed. I didn't have to speculate what this way of life would be like. I had experienced it first hand, and I knew it was good. If I had not had the chance to experience it, if it had been up to me to imagine what it would be like, it would have been hard to be so sure. Maybe I wouldn't have had the certainty that it was worth striving for.

LXF But how do you provide a similar enthusiasm to someone who hasn't had your experience?
RMS But nowadays you can all have this experience. The free software community today is the same kind of community as it was before. But now there are a hundred thousand people contributing. Anyone can get hold of an old PC for little money, and have this experience.

What we can do is talk to people more about freedom. Point out to them that this isn't just a matter of convenient, powerful, reliable software. Yes, GNU/Linux is powerful, yes it's reliable, yes it's cool. Yes you can get it at a low price if you want, but there is something more important there. It allows you to use a computer in freedom.

LXF You have often said that if the options were proprietary software or nothing at all, you would rather have nothing at all. But part of this passion, and the fact that you understand this freedom, comes from the fact that you are a hacker. How can I explain this enthusiasm to someone who isn't involved so heavily in computers. How can I explain it to my mother?
RMS First of all you explain that programs are like recipes. Your mother has probably shared her recipes with friends. And she has probably shared theirs. Imagine if she was told that she may not change that recipe. Imagine if she was told that she wasn't allowed to share that recipe with her friends. This is a very good analogy that helps people to understand.

"It would give people a chance to live in freedom"

LXF But the difference lies in that everybody is capable of cracking an egg, but not everybody is capable of reprogramming a piece of code.
RMS But the point is that the analogy should help people to understand, and if they've ever used a computer program that frustrates them and they have been unable to change it, then they can begin to understand. They may begin to assume that this is no use to them because they personally cannot change it. You have to show them that they can pay somebody else to change it for them, or they can go to a family member and ask them to change it, or they can go seduce a programmer and ask them to change it for them. You become a part of a community, and you soon realise that people can and will do things for each other.

The community is one of the key elements for Stallman's ideologies to be successful. There has to be a sense of belonging – and a real contact with those who people are putting the effort in for. When there isn't an enforced financial transaction for your efforts, there is an increased necessity for praise and feedback from those who are benefiting from the product. And as a reasonably recent convert to all things GNU/Linux, the atypical friendliness of this community is striking.

LXF It is quite surprising once you are in this community, quite how willing people are to »

Free history

Stallman got his first taste of the free software ideal in 1971 when he joined the hackers at MIT Artificial Intelligence Lab. This community collapsed in the face of technological and political change during the early 80s.

At this point, Stallman says he faced "a stark moral choice." He could join the proprietary software machine and propagate the idea that ideas were commodities, he could leave the software community altogether, or he could attempt to change the way people regarded software.

In January 1984, Stallman quit his job at MIT and set about creating his Gnu's Not Unix (GNU) operating system and applications. This was followed by the creation of the Free Software Foundation in 1985.

In 1991, Linus Torvalds released his Unix compatible kernel and this, combined with much of the nascent GNU Hurd project resulted in what we now regard as Linux, or more properly, GNU/Linux.

help. People can go out of their way to improve things, and help you out if you request it. The willingness to be told about bugs and to help solve problems is remarkable.

RMS Absolutely. This is a thing that many people don't understand. People have a tendency to complain about bugs in a program to everybody except the one person that they should be telling, which is the developer of the program and the person who may be able to do something about it. And he is the person who needs to know. And yet people find all sorts of ways to avoid thinking like this. So in the GNU project we have a policy: when somebody reports a bug, unless he is totally callous and stupid about it, you thank him, because he is doing something helpful. Without bug reports from the users, you cannot get it right. Without feedback, you cannot learn. So reporting bugs is part of the users responsibility. It's the way they contribute to our community, the way they help to build it. So we thank them. "You showed me this fault in my program, which was there already, but because you showed it to me I can now fix it."

LXF And hopefully, when one is reporting a bug in free software, one reacts differently than perhaps they would when reporting a bug in proprietary software. If I have paid hundreds of pounds for an application, and it is riddled with bugs, then I am going to be rightly cross. But if I have acquired a piece of software freely, made by someone for the betterment of the community, then I am hopefully going to be more forgiving in my reaction to problems. To be all tetchy and cross about it, you'd have to be a bit evil.

RMS Well, there are people who are like that, and of course it's not fun when they do it, but as producers of software we have to learn that the reporting of a bug is a very important part of our community. The users have to learn this so that they will report bugs to the software producers. For example, if you had bad breath, and I smelt it, would you want me to tell *you*, or would you want me to tell all our friends. That's a harmful way to treat someone. If I have bad breath, I want you to tell *me*, preferably privately, and then I will fix the problem with a toothbrush and toothpaste. And then nobody else would have to smell it. Whereas if you told my friends, and didn't tell me, that would be the nastiest possible thing you could do. Think of that whenever you find a bug in a free software application.

LXF It all seems very friendly and helpful for such a dangerous and radical group as you are so often portrayed.

RMS People often make an assumption that freedom is the opposite of responsibility. That's not true here. Here you have freedom and responsibility on one side, and being dominated and divided on the other side. Remember we are talking about the freedom to co-operate with other people, the freedom to be a responsible person in a socially responsible way, and contribute to your community. The Free Software movement and community encourage people to help each other. We don't just say: "you should have the freedom to tell other people to go to hell", we say: "you should have freedom to help other people." We encourage people to work together and co-operate.

But within this community there are, of course divisions. The most significant of which is the split between the Free Software community, and the Open Source community. Such a divide is not at all helpful to the cause, meaning that people who really should have been reaching for the same goal are now working at cross-purposes. But in Stallman's eyes, this is not merely a case of two rival groups working for the same thing – in his view the Open Source community are missing the whole point. One of the key offenders in this split is book publisher Tim O'Reilly. He proclaimed to be very much against the powers of proprietary software, and then began publishing manuals for all manner of non-free software. And more importantly to Stallman, he publishes them as 'non-free'.

LXF So you think that Tim O'Reilly is going about things the wrong way then?

RMS Absolutely. O'Reilly presents himself as a supporter of our community, while mostly refusing to contribute to it.

By publishing most of those manuals as non-free he is acting against the interests of our community. We need to have free documentation to go with our free software. There are many programs that don't have free documentation, because O'Reilly has convinced program writers to publish non-free.

He is very clever, and very good at subtly shifting the service.

So why was there the splitting of the community. The term 'Open Source' was originally invented to be a kind of dis-ambiguator of the word 'free'. Well, that was one of the goals. But another of the goals was to avoid using a phrase less likely to make business uncomfortable. And of course there was Eric Raymond, one of the most vocal supporters of the Open Source movement, and probably the man who has been more personally abusive to Stallman than anyone else working within free software.

LXF Open Source, and its more business-friendly appearance was something that Eric Raymond encouraged, so what was his angle?

RMS He simply doesn't agree that people are entitled to the freedom to share and change software. He likes having that freedom, but he rejects the idea that people should be entitled to it. So he rejects the idea that non-free software is wrong.

"Without feedback you cannot learn"

He wrote an article suggesting that Open Source is merely a very clever move to con big business into making free software without realising it.

He proclaims to agree with all my principles, but he is trying to mislead people about where I stand. He is trying to spread his views. I am more or less an articulate spokes-person for my views. If he can give people the impression that I really stand where he does, it will boost his views.

I feel it's somewhat disrespectful that he attempts to misrepresent my views.

LXF There was certainly a divide between the two communities, and they didn't appear to want you to be a part of it.

RMS They coined this term [Open Source], and I saw that it was going to be used to separate our community's activity from the philosophy that we have in the GNU project, and I realised that I'd better not go along with that. They wanted me to not make a fuss and to allow them to take the lead as to what philosophy to present to people. They have a right to have and promote their views, but they cannot expect me to support them.

LXF There is undoubtedly confusion about the use of the term 'free'. A Swedish speaker at the CODE conference gave a talk on why the dot.com industry was doing so badly because of the approach that everything should be free. He had unfortunately completely misunderstood the use of the word 'free' due to the language barrier and it's numerous interpretations. This ambiguity cannot be a positive thing?

RMS There is no good word to use in English.

LXF Can we invent one?

RMS No! No we can't. In theory it would be possible to invent one, but to do so and to switch to it would not work. In a different situation, if people all wanted to do it, then we could. But in the current situation, where the Open Source movement are working so hard to take people away from the Free Software movement, to make such a change would be to lose the share of the attention that we've still got.

LXF The speaker went on to say that the attitude was causing people to make no money and to fail.

RMS Well, the worst that can happen is that they can get new jobs. They're not going to starve to death. But those are secondary issues. How much money programmers make, or how many people can have jobs as programmers, is not as important as the freedom of all computer users.

LXF Do you feel that as a result of the fracturing of the two movements, that they have the better term?

RMS Well, yes. But there is nothing that can be done about it. In fact, there is no good term for this in English. Open Source wouldn't be good for us. It's very effective for them because it avoids raising issues. 'Open' is a warm fuzzy term. It sounds like there's something good about it, and it is very vague about what it is.

LXF Would it be possible to perform some sort of linguistic appropriation, perhaps in the same way as homosexual groups have taken the word 'queer' and completely appropriated it, completely subverting it. Can you subvert 'Open Source'?

RMS No. For one thing it would be disrespectful. They have a right to promote their views, and a right to say what they stand for, make up a name for themselves, and use it. And the other reason is that there are so many companies that use the term, and agree with they way they use it.

There are too many people standing by the use of the term in its current meaning, and by comparison we would be insignificant. We would simply lose people's respect because we would be acting in a less upright fashion.

He remains this honourable always. Despite his erratic behaviour and temper, this is a man of his words. He won't be drawn into slagging people off, not even those who make a living out of mocking him. He will tell you the truth, as he sees, and lives, it. It's actually rather unnerving, but I think only because it is quite so unusual.

His powerful principles lead him to some strong opinions. One area he is very opinionated in is patenting. They are of critical relevance to the success of the Free Software Foundation.

LXF You are very passionate about the misuse of patents. When and where, if ever, do you think they're appropriate?

RMS Patents should simply be excluded from software and anything else that isn't specifically physical. Patents affect different fields differently, and there is no fundamental reason why there should be a uniform patent for all fields. But what I can see is that software is at one extreme, and say at the other extreme is pharmaceuticals. The amount of expense it takes to develop something of any given complexity in software is very low, and in pharmaceuticals is very high, and everything else is in between them. Software is so low because you are dealing with mathematical entities, abstractions, and any in any other kind of engineering you are dealing with matter. And matter is perverse.

Since the intelligence of people working in both these fields is basically the same, if you're working on something hard, you do smaller systems, and if you are working on something simple you do more complex systems. That's why software packages are so much more complicated than mechanical and technical designs typically are. So the result is, people develop very large software packages, and because of this, a lot of ideas are used in it. A word processor has a lot of ideas in it, if it is to be competitive. But a lot of those ideas may already be patented, and if they are, how are you going to develop a competitive word processor? If you want to develop a better one, you will probably have a few new ideas, but you have to use lots of well known ideas that have already been used by others. And if those ideas are patented you can't develop your product, even if it is an interesting and superior word processor.

It is explicit in the US constitution that the purpose of patents is to promote progress. It is the same in the UK. They are not an entitlement, they are not a natural right for inventors. They are an artificial government system to promote progress. So the question is, does it promote progress? In software, patents obstruct progress. I'm not against patents in a field of manufacture, but I am not an expert on this. But when you are talking about big companies, they don't have any inherent rights. Only people do. Businesses cannot complain about injustice. They try to claim this right, which is why it is so important to reject these claims whenever they appear.

Software is not just a concern for large companies, because software patents restrict individuals in the use of their own computers. They restrict individuals from working together in groups. This means that the effects of patents in software is very different from that in, for instance, the automobile industry.

LXF People claim that free software prevents people from making money.

RMS I don't care. I don't care. Free software respects our freedom. It is important because it enables us to use computers and have freedom.

"Patents should be excluded from software"

And so it always comes back to this. This is the key statement for Stallman. No matter what the subject, if it gets in the way of freedom,it is irrelevant. And it is undoubtedly this central message that has enabled him to be driven so hard, and to have achieved so much.

LXF Finally, is there anything to you that is more important than the principle of free software?

RMS Well, the survival of humanity I suppose, but in the world of software, no. LXF

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**

Any more for Tibet

I have full sets of *LXF* and *LM* since the first issues, including all the cover CDs. More use to the Tibetans than to me.

Who/where should I send them to?

Guy Inchbald, *via email*

›› Thanks for your generosity! Anyone else who is interested in helping out the Tibetans we mentioned last issue should send stuff in this direction:

FAO Mr Ugyen Kyab,
Institute of Buddhist Dialects Computer Centre,
P.O. McLeod Ganj,
Dharamsala,
Himachal Pradesh176219
India

So far we have sent them some general Linux books, copies of *Linux Format* (the full set now thanks to Guy) and latest distros of Mandrake, SuSE and RedHat.

If you can think of a group or project who could genuinely use some help, write in and let us know – at the very least we can send off a small parcel of Linux related goodies.

New look

Just thought I'd let you know that I really think you've done a good job on the new look *Linux Format*. I disliked the early covers and thought the magazine was rather thin for the money. I like the new direction, you've made it into a serious magazine (and no more bloody penguins, thanks). I also think that your spelling has improved, the usual mistakes with "its" and It's" are not so visible. I'll be subscribing!

Neil Lucock, *via email*

★Letter of the month

This month's winner receives a copy of **The New XFree86**

Credit where it's due

I was recently up on the Red Hat website and saw an advert for a 'charity' credit card sponsoring the Linux Fund issued by MBNA America. As an MBNA Europe customer, I was very interested and rang the company in the hope that they would run a similar scheme. They don't. The advert features a Credit Card with a Tux logo and offers a 6-inch plush Tux to boot. What I want to know is if this Linux Fund will result in money going to worthwhile development and if it does if there is some way that you might be able to help getting MBNA to set up a similar scheme in Europe.

If I thought that this card was helping Linux development, I would use it in preference to all my other credit cards. I want Linux to become a viable alternative to Windows.

I enjoy reading your magazine very much, but find all the different Linux distros a bit confusing. Would it be possible to run a poll to find a preferred distro and then run more articles on running and configuring that distro specifically?

Roger Branwhite, *via email*

›› The good news on the credit card front is that a Linux Fund card is expected to be introduced to the UK and Republic of Ireland this year, and you can be sure we'll let you know when this happens.

One reported downside of the card in the USA is that there is no statement telling you just how much has been contributed to the cause, nor have we been able to get any information on how the contributions are worked out. Hopefully these details will be clearer by the time the card is introduced to the UK market.

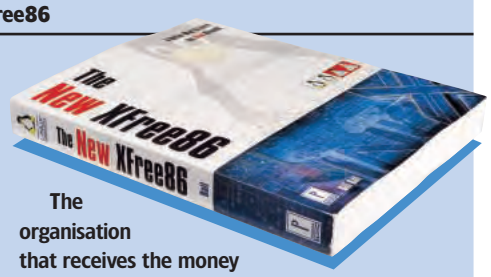

The organisation that receives the money collected is Linux Fund (www.linuxfund.org), who use it to give grants to developers of open source projects (to help with equipment and so on) and also offer limited scholarships to promising developers.

The type of projects funded are usually ones which will be a long term benefit to the Linux community, but have little commercial potential and therefore find it difficult to get alternative funding.

In terms of popular distros, SuSE, Red Hat and Mandrake lead the pack. We have run articles on installing all of these to coincide with coverdiscs. The trouble with doing distro specific articles is that there are many times more people that aren't using any given distribution than are. Added to that the fact that people want to use Linux for different things, have different hardware etc, and you have a piece that's going to run for quite a few pages. If there are specific areas of installation or configuration you're having trouble with, you can always email us or pop on to the forums, or even check out the online documentation from the distro developers.

This month you would have received Bill Ball's *The New Xfree86* (Prima Tech, ISBN 0-76153-152-1), but since you chose not to supply even the vaguest of addresses, we're going to denote it to a worthy cause instead...

›› We did obviously try to create a magazine that reflected what most people wanted - I'm glad that in your case we've got it right. I like Tux, but it was getting a bit tedious seeing him all over the place!

Disappointingly we haven't received a single letter complaining about the changes we've made in the magazine. Come on, what's wrong with you lot? Surely someone must have a picky complaint or two!

More GIMP

Is there any chance that some stable *GIMP* plug-ins can be put on one of your future CD's?

I, for one, am not yet on the web (how Neanderthal can you get?). I guess there are many *GIMP* users out there who would appreciate these additional plug-ins.

Also I would like to see some tutorials on *GIMP* mastery... Any chance?

Herb Spencer, *via a friends email*

›› Oh, okay then. We have run *GIMP* tutorials in the past, but we'll probably put together a *GIMP* special feature soon, with lots of hints and short tutorials. Will that do you?

DVD gripe

Just read about the DVD last night. Guess what? Couldn't get a copy from Future Publishing. They say they've sold out. How about a re-release, pleeeeaaassse!? If you produce any more DVDs, can we have advance notice, so that people can subscribe? If so, count me in for copies!

Thanks for a great mag. Keep up the good work.

Richard Broughton, *Lincoln*

DVD gripe 2

I am a regular subscriber to your magazine and I guess like most people who have a subscription to *LXF*, I can at least install Linux on my machine. You will understand then that SuSE 7.1 Live on the April CD was not exactly top of my wish list. That said I must say I don't only subscribe because I enjoy the magazine but also as one small way of supporting Linux, so when the CD contains good advocacy material, I can only support you. Not so with the April issue though because I got to page 98 and started to read about all the good things all us subscribers missed on the DVD. If you wanted to *really* annoy your subscribers you could not have done a better job. This may sound like a moan and if it does, I'm sorry but I thought I would point out this little error so you can avoid it happening again and give you a chance to apologise to us loyal readers. In every other way I think the mag is getting better with every issue.

Tom Ormes, *via email*

›› Firstly, let me apologise to all the disappointed readers who were unable to obtain a copy of the DVD. This is basically because we just underestimated the demand for the DVD – it seems Linux users are much more keen on them than the average Wintel devotee. There were other problems associated with distribution which didn't help, but the fact that we ran out of the reserve copies so quickly is our fault.

I'm sure it was very frustrating for a great number of you, and I can only apologise profusely and say that collectively we have learnt our lesson. We will be producing DVDs in the future, but we are looking at a number of ways the process can be managed more effectively. One thing I can definitely promise you is that subscribers will find out well in advance, and we will find a way to guarantee you a copy of the DVD if you want it.

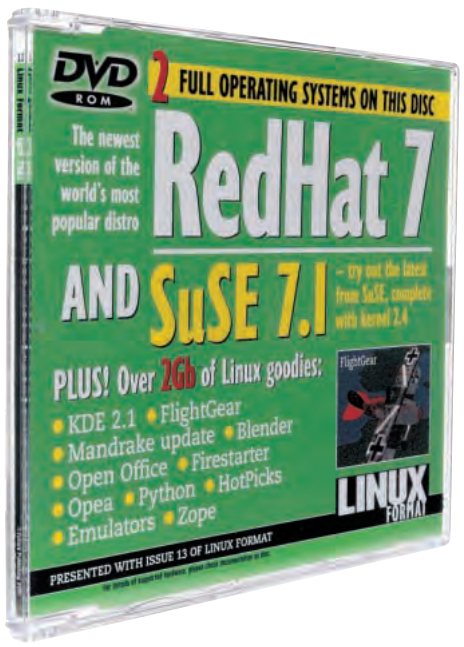

The LXF DVD: it's rarer than hen's teeth.

Where are the tools?

With reference to April '01 edition page 45 Apache Toolbox, this site cannot be located from my browser. Is this a fire wall problem or is the URL inaccurate?

Alan Spandrzyk, *via email*

›› Sorry about that, a few key letters went astray there. The correct address for anybody having trouble is www.apachetoolbox.com

CD con?

I'm fed up, I bought your magazine and got conned again. OK so I can't read, but believed I was getting a full version of SuSE7.1 on a DVD. OK it is brilliant and clever, and the *GIMP* just keeps getting better.

But what am I to do? If I buy the boxed version it's out of date, and beyond my means to upgrade, well nearly everything.

Please, please, please could you arrange to have upgrades for SuSE (and/or Mandrake etc) with the trial version, then I could buy the box and have the upgrade ready to use.

Even better, just offer the option of a full install!

Jonathan Chetwynd, *via email*

›› The SuSE Live 7.1 CD was designed (by SuSE, not by us) to run directly from the CD. The benefit of this is that users can try out the whole SuSE 7.1 experience without having to do a full install (usually over the top of an already working distro), and also making it possible to try out SuSE Linux without having to repartition your drive first.

Getting this to work is quite technically demanding, and also takes up a fair amount of the CD space. Including a standard install set too would have taken up much more of the CD, and there would have been less room for the packages that comprise SuSE 7.1 – giving you an incomplete product.

We thought the benefit of being able to try out this distribution without having to install it was one that would be appreciated by a significant number of readers. I'm sorry if you weren't one of them.

As for putting package updates on CD, we could easily fill a whole CD with updates for Mandrake and SuSE every month, never mind Red Hat. However, whenever we feature things on the CD, e.g. *KDE*, we try to put as many package combinations on as possible, to make it easier for users of all distros. The common denominator is a source tarball, which can be used by everyone running Linux, including those using non-x86 architectures. We try to be inclusive, and choose the packages that will be most use, but we can't fit everything on the CD!

Networking

A quick request from someone who is fairly new to both Linux and networking – how about an article in your magazine about setting up a home network with Linux? I'm sure I am not the only reader with ambitious ideas and not enough knowledge!

Also, do you now if any of the new wireless home networking kits are compatible with Linux? Most of what I have seen is aimed at Windows users.

Nick Baird, *Gloucester*

›› Networking has been duly added to our list of things to cover. There is patchy support for wireless networking devices but, as ever, there is a lack of standardisation here making it difficult to produce comprehensive solutions. Hopefully we'll be able to pull together the current state of play in the Networking piece.

> **"I started to read about all the things we missed on the DVD. If you wanted to really annoy your subscribers you could not have done a better job!"**

Read the (fine) manual!

Oh dear: another issue of *Linux Format* and another reader complaining of the difficulty of installing software and the steep learning curve of Linux.

But Linux is a very powerful operating system ››

« and there is simply no point in trying to use it if you're not going to take the time to read books and manuals and literally struggle to get it working. I wonder how many readers would still feel comfortable if they were presented with a Linux box with no X-display and told to work from the command line?

I realise that this will probably offend a lot of people and I also realise that it is not their fault that they approach Linux with the wrong attitude: how many people actually read the manual for their car radio? How many sit down and read right through the video machine manual? We live in a society where there is no time allocated for learning and reading and experimenting with new technology. Thus people take the same attitude with Linux.

"Oh dear: another issue of Linux Format and another reader complaining of the difficulty of installing software and the steep learning curve of Linux."

Linux is a very powerful operating system and to get the best out of it you really do need to sit down and actually read about it: that does not simply involve reading *Linux Format* every month: good though it is, it can not in any way be a replacement for reading man pages, info pages, howtos and books in general on Linux. Back in the days of MS DOS and *Wordperfect 5.1* we all had to learn those few commands to get us by inside DOS. Linux is much the same: although X is very helpful and productive, it is a necessary skill to be able to survive for a console: the number of times I've had to repair broken X configurations from a console...

I deliberately don't use *KDE* or *GNOME*, though I have both installed as I use applications that depend on them. My window manager is *uwm* which is like nothing on earth and about as far away from Windows as you can get (is an article on *uwm* possible please?). Constantly I find that console programs are more powerful and reliable than their X equivalents: I use *Mutt*, *micq* and *bitchx*: all are very powerful, but you have to read the manual before you're going to get anywhere with them.

So please, let's not have any more of "Please can you make Linux easier to use" because it's not going to be productive to make it easier: you'll end up with Linux having to try and second guess what you really wanted to happen and then you'll end up with a bloated OS with hundreds of 'wizards' which is exactly what Linux is trying to avoid.

Besides, is it really too much to have to learn five installation commands that work for probably 80% of the packages you're going to have to compile and install yourself:

bash$> tar xzvf package.tgz

(Explanation: x = extract from the archive;

z = run through gunzip first to uncompress the archive;

v = verbose – list the files being extracted;

f = filename – the file you want to extract from)

bash$> cd package

bash$package> ./configure

bash$package> make

bash$package> make install

(Explanation: these last three commands are part of the GNU autoconf tools that make creating a package much easier: ./configure grabs information from your system that allows it to generate suitable "make" files. These files are then used by "make" to compile the package, and then "make install" puts the compiled files in the right place)

In fact, the distributions that most people think are the hardest have got some of the most user friendly commands available: the Debian package manager will autosolve dependencies: no wasting hours searching for an rpm. Debian also allows you to compile a new kernel into a package which you then install just like any other package. But none of these helpful tools are of any use if you're not going to take the time to learn how to use them.

So then, book recommendations: *Running Linux* from O'Reilly is a good starting book, while *Linux in a Nutshell* is a very good quick guide to most of the required shell commands, and then there are mailing lists and howtos. At first it takes some time to find out where to look for the answers. But once you've found where the answers are, you can keep going back for more.

Good luck, but before complaining, try hunting down the answers yourself. People on mailing lists are far more likely to want to help you solve your problem if you've done the groundwork already.

Matthew Sackman, *Nottingham,*

We covered performance tips in issue 7.

» Controversial maybe, but does Matthew have a point? Are we living in an appliance society where we expect computers to interpret what we want to do rather than allow us to make specific requests? Creating an operating system that is easy to use for beginners, but gives power to those who know how to use it is a problem people have been trying to crack since the days of punch cards.

Is this an insoluble problem, or is Linux flexible enough to allow both approaches? Write in and let us know your opinions!

Trimming Linux

In common, I suspect, with many other users, my Linux install is full of stuff which I never use and whose purpose I am unsure of. So how about an article telling us what we can get rid of?

Window managers/GUIs, browsers, e-mailers, newsreaders, text editors: we can all work out for ourselves what we're not using and un-install accordingly. But a small push in the right direction wouldn't hurt.

Programming: a bit trickier. I know we need *gcc* to re-compile the kernel and to install a program from source, but what about the other C-related stuff? Does *gcc* need it, or is it just there in case we feel like developing our own apps? Java: we need Java/Javascript for the full web experience, but how much of it? Perl, tcl/tk, etc: do we need these at all if we're not programming them?

Apache and *Sendmail*: my Windows partition seems to manage without a web server and a mail server, so why do they keep popping up in Linux? I've seen more than one newsgroup post asking "why is my Linux so slow?" being answered with "Remove *Apache*".

This issue could of course be addressed the other way: 'Buried Treasure'. What really useful apps are lurking in the average installation that the average user would not be aware of ?

Apologies if you've addressed all this in a copy of *LXF* which I didn't buy.

David Cowie, *via email*

» You'd probably be better off not installing the things you don't want in the first place, as uninstalling them later can be slightly trickier. Many of the services that are started normally on a Linux box may be of no use to you if you use it strictly as a home desktop machine, but most of these services, particularly the Internet ones, are only started when you use them, so do not waste resources. You can of course choose to shut down certain services. *Apache* isn't really necessary unless you are using your machine as a server or web development. Although it may be difficult to tell you what things to turn off, I guess a feature which actually tells you what these things do might be useful, then you could make up your own mind. We did do a feature which touched on this back in issue 7 (how could you possibly miss a copy?).

Inevitably...

The Sega SC-3000 was the only computer ever manufactured by Sega. It had a 4MHz processor, 18k ROM, 8k RAM and 16k Virtual RAM. It supported 16 colours, each with 15 shades. The machine came boxed with games and a tape drive. A BASIC cartridge was available as well, but this took up a great deal of the machine's memory. The Japanese version of the machine was white.

I just noticed that you didn't know what that was in the article about emulators.

Also, I noticed that there seems to be no Linux Users Group in Glasgow. Anyone who wants to talk Linux will normally find me at the Glasgow 2600 meeting which takes place at the payphones on platform one at Central Station.

Chaotic, *via email.*

›› Thanks for the info. We did actually know, but the poor, tired individual who was captioning stuff late one night was obviously suffering from sleep deprivation.

Hmm, thanks for the LUG info too, we'll be sure to drop in next time we're in the area.

The 'Buried Treasure' idea is sort of addressed by our regular What on Earth feature, though this has expanded to include technologies like SOAP, for instance. Thanks for the ideas anyway!

New forums

I enjoyed your old-style web-forum much better as it was easier to navigate and easier to see what was happening!

Could you perhaps reconsider?

Jim, *via email*

›› Possibly. The trouble with the old forum was it became unstable in the heavy traffic situations which often occurred on the website. So while we're open to suggestion on alternatives, we're definitely not going back to the old one.

Disappointed

As you have introduced a DVD issue as an alternative to the CD-ROM, it is disappointing as a CD-ROM user to not get those programs described in the DVD-pages section of your mag.

And you are not supplying the applications reviewed in HotPicks and Roundup sections each issue either.

In my view it would be a better solution to take a leaf out of *PC Plus's* book and supply two CR-ROMS in a single case. Then you could supply, perhaps, a distribution on the first CD with HotPicks, Roundup, games and the other stuff every month. Skip the DVD (no joke) and supply a pair of packed CD-ROMs in a single case. Then everybody can benefit from the same selection of distros and programs. Hope to see FlighthGear on a CD-ROM soon.

Geir Hungnes, *Norway geirh@postkassa.no*

›› Thanks for the suggestion. The problem with it, as you might expect is money. It would simply cost far more to do two CDs. Normally we do try and get related stuff from the magazine on the disc, but this was not possible with the SuSE 7.1 CD, because of the way the Live version works. As for your Flight Gear request, I'm afraid all the textures and backgrounds we provided, as well as the source itself, would probably fill more than two standard CDs.

Converted!

Having progressed through DOS x, Windows 3 x and 9 x (and back again on numerous occasions), as well as various versions of *Office* and its constituent parts, I feel myself to be an advanced Microsoft user able to troubleshoot most problems that arise. However, I never liked the idea that so much of the essential software on my system – bar utilities and anti-virus applications – was produced by one company and so when I set about building my own PC I was determined to use software from alternative sources; *Lotus Smartsuite* instead of Microsoft *Office*, *Paradox* instead of *Access*, *Opera* instead of *IE*, *Visual Page* instead of *FrontPage* – you get the picture.

Unfortunately, there seemed to be no alternative to Windows that did not involve too steep a learning curve. I experimented with BeOS, and Red Hat 5.2 but was not won over. When I saw SuSE 7.1 Professional advertised, read some complimentary reviews as well as tried out your cover CD I decided to have another look at Linux.

I was impressed right from the start. The box contained comprehensive documentation and numerous CDs (and a DVD) and the setup was a joy. Having put my PC together myself I was anticipating possible incompatibilities but the setup and installation were very smooth, albeit time-consuming. I haven't yet needed to call upon SuSE technical support although I did do a little surfing to find how I could set up my Epson 680 on Linux. My older printers, an Epson 600 and LX400, both work without a hitch.

Setting up the modem and Internet connection was a breeze and my surfing seems to be twice as fast on the Netscape version supplied. However, my first download was *Opera* for Linux which is my browser of choice at the moment – though *Konqueror* comes a close second.The CDR/W works great although I am still in search of more comprehensive software for it.

KOffice and *StarOffice* are more than adequate for all the business type applications I require and Quanta is my HTML editor until *1st Page 2000* is ported to Linux (one day I hope). I find both *GNOME* and *KDE* easy to use and powerful and after two weeks of continued use – including unavoidable tinkering ;-), I think my troubleshooting days may at last be over. Not one Microsoft Blue Screen of Death encountered.

When WINE or similar emulator becomes more effective I will no longer be dual booting!

Steve Payne, *via email*

> **"The trouble with the old forum was it became unstable in the heavy traffic situations which often occurred on the website. We're not going back."**

Mailserver Hot Topics

We have introduced Mailserver Hot Topics to help gauge your opinions on the things that matter most. Please feel free to continue writing in on any subject you like (except the glues we use on the CDs, everyone's heard enough about that), but we would be extra keen to hear your views on the hot topic of the moment.Without further ado, next month's topic is: What software applications does Linux desperately need?

Submission advice

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

WRITE TO US AT:
Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email:
linuxformat@futurenet.co.uk

User Groups

Wherever you are in the world, there'll be a Linux User Group somewhere near you. There are hundreds of dedicated User Groups all over the UK alone, so find the one nearest to you now! And if you're outside the UK, you needn't feel left out either, 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/links.php or at www.lug.org.co.uk

1 Hampshire
URL www.hants.lug.org.uk
Contact Ken Adams

2 Bristol & Bath
URL www.bristol.lug.org.uk
Contact Dave D

3 Scottish
URL www.scottish.lug.org.uk
Contact Tony Dyer

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

5 Bromcom (Kent)
REVISED DETAILS
URL www.kent.lug.org.uk
Contact John Mills

6 Brighton
URL www.brighton.lug.org.uk
Contact Johnathan Swan

7 Sussex
URL www.sussex.lug.org.uk
Contact Mike Pedley

8 Northants
URL www.northhants.lug.org.uk
Contact Kevin Taylor

9 Anglian
URL www.anglian.lug.org.uk
Contact Martyn Drake

10 Milton Keynes
URL www.mk.lug.org.uk
Contact Denny De La Haye

11 Doncaster
URL www.doncaster.lug.org.uk
Contact Andy Smith

12 South East
URL www.southeast.lug.org.uk
Contact Ian Reason

13 West Wales
URL www.west-wales.lug.org.uk
Contact Ken Adams

14 Wolves
URL www.wolves.lug.org.uk
Contact Jono Bacon

15 Peterborough
URL www.peterboro.lug.org.uk
Contact Steve Gallagher

16 Edinburgh
URL www.edinburgh.lug.org.uk
Contact Alistair Murray

17 Tyneside
URL www.tyneside.lug.org.uk
Contact Brian Ronald

18 Leicester
URL www.leicester.lug.org.uk
Contact Clive Jones

19 Greater London
URL gllug.linux.co.uk

20 Surrey
URL www.surrey.lug.org.uk
Contact James Chivers

21 Cambridge
URL www.cam-lug.org

22 Devon & Cornwall
URL www.lug.termisoc.org
Contact Alex Charrett

23 Essex
URL www.epos.demon.co.uk

24 Manchester
URL www.manlug.mcc.ac.uk
Contact Ted Harding

25 West Yorkshire
URL www.scs.leeds.ac.uk/wylug

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

29 North East
URL www.shofar.freeserve.co.uk/NELUG

30 London
URL www.lonix.org.uk

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

40 North Wales
URL www.northwales.lug.org.uk
Contact Jonathan Cole

41 Midlands
»NEW GROUP
URL www.midlandsLUG.cjb.net
Contact Pete Thompson

42 South Cumbria
»NEW GROUP
URL clug.lug.org.uk
Contact Jamie Dainton

43 Dorset
»NEW GROUP
URL www.dorset.lug.org.uk
Contact Beanz and Tracy

44 Hertfordshire
»NEW GROUP
URL www.herts.lug.org.uk
Contact Nicolas Pike

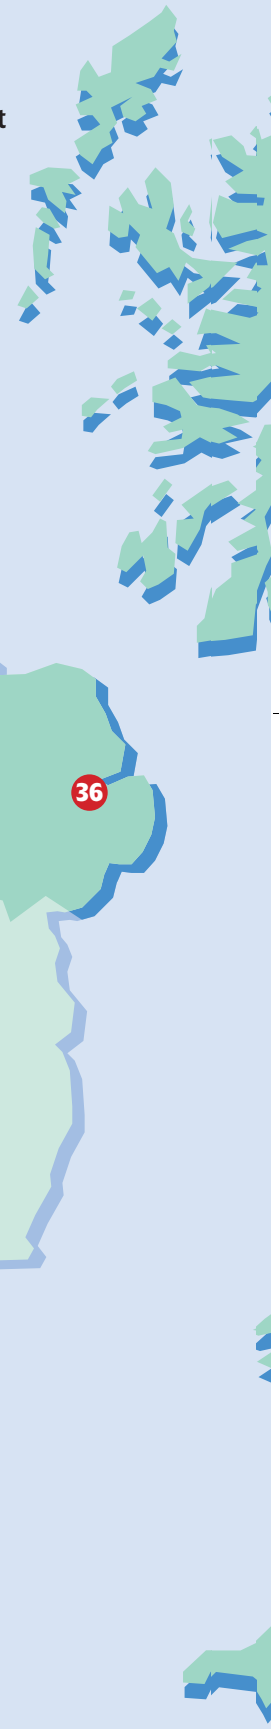

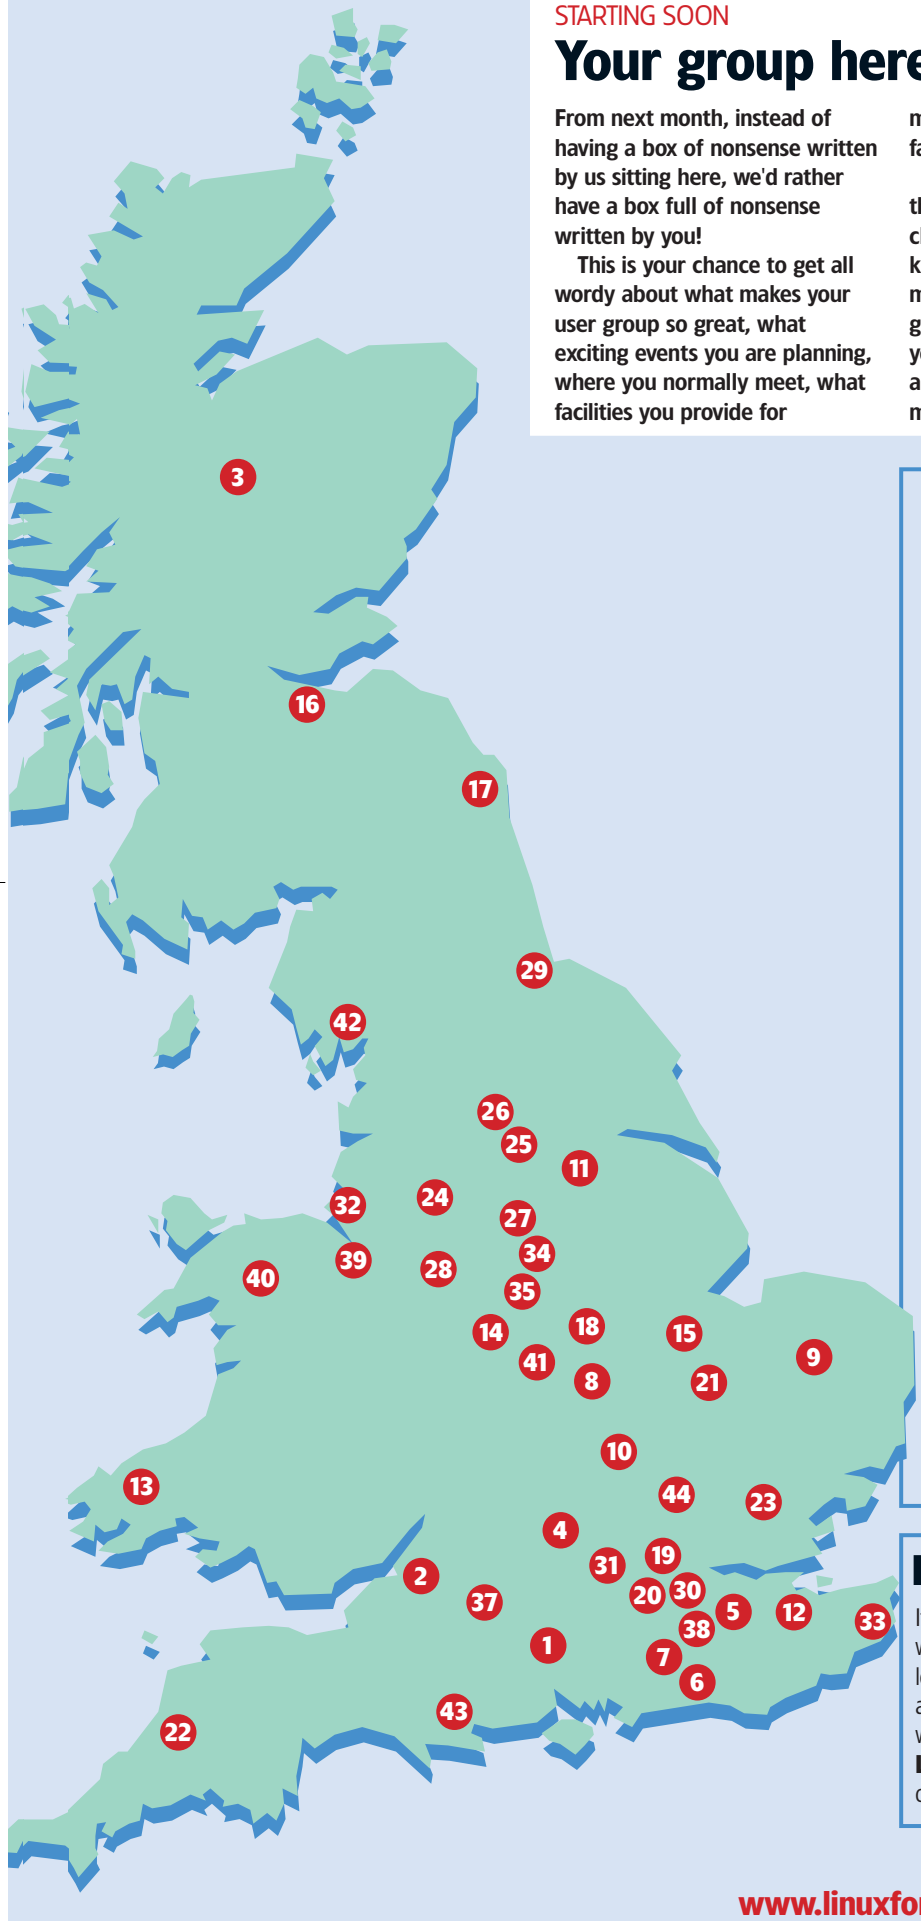

STARTING SOON

Your group here.....

From next month, instead of having a box of nonsense written by us sitting here, we'd rather have a box full of nonsense written by you!

This is your chance to get all wordy about what makes your user group so great, what exciting events you are planning, where you normally meet, what facilities you provide for members and any other claims to fame you can think of.

Basically, we're offering you the opportunity to promote your club a bit more – you never know, you may attract a few more members or celebrity guests. Include any information you think might be relevant, and a nice photo of one of your meetings/events if possible, and we'll propel you to fame on these very pages.

You can email us your submission at linuxformat@futurenet.co.uk (please use the subject LUG details or something similar), or alternatively by post to: LUG info, Linux Format, 30 Monmouth Street, Bath BA1 2BW. Hope to hear from you soon!

Worldwide Linux User Groups

Africa

STELLENBOSCH
URL www.entropy.sun.ac.za/
Email ixion@entropy.sun.ac.za

Australia

ADELAIDE LUG
URL www.linuxsa.org.au
Email mtippet@anu.edu.au

MELBOURNE, VICTORIA
URL www.luv.asn.au
Contact luv-committee @luv.asn.au.

PERTH
URL plug.linux. org.au

Europe

EIRE
URL www.linux.ie
Email root@linux.ie

URL www.dilu.org
Email glossary@dilu.org

GHENT
URL llsgg.rug.ac.be/
Email wvdputte@llsgg.rug.ac.be

GOTHENBURG
URL nain.oso.chalmers.se/LUGG/index.html

LISBON
URL
www.students.iscte.pt/~a12593/gul.html
Email Paulo.Trezentos@iscte.pt

AUVERGNE
URL www.linux-arverne.org/
Email Cyril.Hansen@wanadoo.fr

India

URL www.river-valley.com/tux/index.html/
Email anil@river-valley.com

URL www.linux-india.org
Email newsmaster@linux-india.org

North America

ALASKA
URL www.aklug.org/index.html
Email deem@wdm.com

CLARKSVILLE, TN
URL http://www.cllug.org
IRC irc.midsouth.net #Linux
Email tux@cllug.org

LOS ANGELES
URL www.lalugs.org/
Email dank@alumni.caltech.edu

BAY AREA
URL www.balug.org/
Email aftyde@balug.org

DENVER
URL spot.elfwerks.com/~clue/
Email: lynnd@ihs.com

TAMPA
URL terrym.com/slug/index.html
Email paulf@quillandmouse.com

BATON ROUGE
URL www.brlug.net/
Email dpuryear@usa.net

VIRGINIA TECH
URL corvette.me.vt.edu/pages/ index.html
Email nega@vt.edu

VIRGINIA TECH
Email nclug@nclug.org
Contact Mat Taggart

South America

SAO PAULO
URL gul.linux.ime.usp.br/
Contact gul@ime.usp.br

BUENOS AIRES
Email dcoletti@impost.com.ar

MONTEVIDEO
URL www.linux.org.uy/

LIMA
URL linux.unired.net.pe/
Email linux@unired.net.pe

Linux User Group 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 name of your group, the location, contact name, website address and any other information. Or alternatively fill in the form on our website at www.linuxformat.co.uk/links.php. Post your details to:
LUGS!, Linux Format, 30 Monmouth Street, Bath, BA1 2BW
or email your details to: **linuxformat@futurenet.co.uk**

Answers

If you are really stuck and the HOWTOs yield no good result, why not write in? Our resident experts will answer even your most complicated problems!

»Experts this month

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email (or even drop in on our forums) and it'll all be taken care of.

David Coulson is a sysadmin with VA Linux and OSDN. He's well versed in networking and web matters, amongst others.

Richard Drummond is an experienced programmer who can answer queries on a variety of subjects. A keen Debian user, he's also our resident Java guru.

Nick Veitch is the editor of the magazine, and answers your easy questions! Or indeed anything to do with grub,.Lilo, netatalk, vi...

Mandrake to go!

Q I installed the Mandrake 7.2 distro from your Christmas 2000 issue on my Dell Latitude C600 laptop preinstalled with Windows98SE. I don't have any problems with most of the installation procedures except that I don't know which Display option to use. Can you help me with this?

More power to you and your great magazine!

Marlun M. Valdez

A Laptops are generally non-standard and have cut down or custom chipsets or custom to fit their limited size. Weird hardware notwithstanding, Linux will happily install on the vast majority of laptops if you perform the installation in a text only mode, so as to avoid having X run if the video card is not correctly supported under XFree86. Fortunately, Dell is a particularly popular brand of laptop and the hardware is well tested under Linux, and as such, many people have already installed Linux on that particular model. All of the information, for almost every piece of hardware on it is available from http://www.linux-laptop.net/. It usually involves downloading customised configuration files and doing a little poking around by hand, but it is well documented and explained in an easy to follow fashion. It is best to start with one major feature, such as X or PCMCIA before moving on to more complicated and less important components.

Live and clicking

Q I bought the April edition of *Linux Format* which had SuSE Linux 7.1 Live on the cover CD. Now I know that this isn't the full edition, however I was wondering whether or not under this version you can setup a modem in order to connect to the Internet? All I'm using is a common-or-garden USR 56.6K modem to dial-up. I have tried running the YaST modem connection program in X-Windows, (i.e. the desktop) and it seemed to be fine, however it didn't write/create the file etc/devmodem as it said it had.

I am a newbie with Linux (I have seen the light and I'm not going back!), so any help would be much appreciated.

Kevin

A It is possible to setup various network options from the SuSE Live CD, but it can never be fully configurable when running of a CD, as ther'es nowhere to save the files to!
As you seem to be interested in Linux, it's probably worth getting hold of a proper distribution which you can install, either from a cover CD, or by downloading the ISOs from the Internet if you're a madman with lots of bandwidth and a CD burner.

Apache and PHP

Q I run SuSE7.0, on a PII 400 with 192Mb RAM and 4Gb and 9Gb SCSI drives. I also run *Apache* 1.3.19 installed at /usr/local/apache/ with module=most and shared=max. This set up works well and I host my web page through a permanent 64k pipe.

I tried installing php4.0.4 according to your instructions in issue 7 (November 2000) as I like the look of pages such as sourceforge.net and linuxformat.co.uk and I would like to emulate these designs in my own pages. But every time I install php4.0.4 and restart the *Apache* server I get the following error:

Invalid command 'LoadModule' perhaps mis-spelled of defined by a module not included in the server configuration.

If I hash (#) out the offending line from httpd.conf 'LoadModule env_module libexec/mod_env.so' the error just moves onto the next line and so on.

Dave

"Weird hardware notwithstanding, Linux will happily install on the vast majority of laptops if you perform the installation in a text only mode"

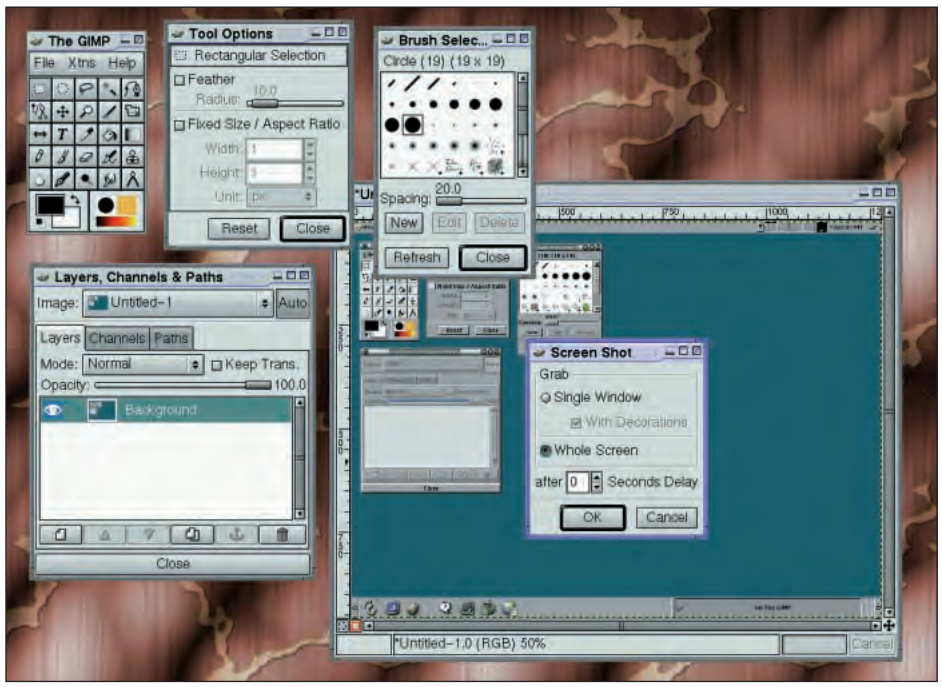

There are a number of ways to grab screen images on your Linux box, including using *GIMP*.

A In order to use PHP as a loadable module, or DSO, in *Apache*, you need to ensure that *Apache* has been compiled with support for shared objects. You can check this by running **httpd -l**, which will list all of the modules compiled into the HTTP daemon. If 'mod_so' is not listed, then you will need to rebuild *Apache* with the '—enable-module=so' option, and any others needed, so that you can use the DSO for PHP with it. You could, of course, compile PHP into Apache, but that would also require Apache to be rebuilt.

Screen printing

Q All I need to know is if this is possible. This afternoon I had to reboot to Windows just to get a couple of screen dumps for something I was writing. It would be a lot easier if there's a GNOME (or whatever) equivalent to the easy PrintScreen|Crl+V sequence in Windows.
James

A Most graphics packages, including *The GIMP* and *Xv* have facilities to take shots of either the whole desktop or individual windows. There is, naturally, a command line tool which does the same thing: **import**, from the ImageMagick package. You can grab the whole desktop with **import -window root desk.jpg** or by selecting the window you want to grab with the **xwininfo** program then using the window ID as the argument for the **-window** switch.

Burning issue

Q OK. I've set up our office server using Linux (Mandrake 7.2 actually) and got the CDR/W working fine, the USB Zip drive working, the USB printer working, samba working, web server etc... But is my boss happy? Is he bobbins! He went into great detail as to how, at home, on his own machine he can use the CDR/W Drive like a hard drive, by preparing the CDR/W Disk in a certain way he could then just drag and drop files to and from it. Now, we have the CDR/W shared through Samba across the network but how can I get it to have this same functionality?
Yakkerty, from the LXF forums

A This sort of burning uses what is known as 'packet writing', which as yet, there is little support for in Linux. You may want to keep an eye on the Packet Writing kernel patches for Linux over at http://packet-cd.sourceforge.net/, but they are, at least at the moment, very much in development and unless you've got plenty of time and patience spare, you're probably best off sticking to burning an ISO created with 'mkisofs'.

Speed daemon

Q I've noticed that every time I look at the running processes on my SuSE 7.1 box there are at least five occurrences of NSCD running. Is this really needed? I'm only using an AMD K6-2 400 and 128Mb RAM and it is running a little slow, so I'm pruning the system down. I haven't gone to the trouble of re-compiling the kernel yet as I think some sensible rc script pruning should have a reasonable effect.
Oliver Thomas

A Nscd is a fairly simple DNS caching daemon, and even if there are five processes running, it's unlikely that it will cause any great slow down on the system. In fact, it is probably best to keep it running, as it will improve the speed at which DNS lookups are done, assuming that it is configured correctly, when a program has to do repeated lookups on a particular hostname. You may want to run 'top' from a terminal when the machine is up and running, so you can see which processes are using the most RAM or CPU time.

Distro dilemma

Q I have tried to find a reliable distro, but have failed to find one without many problems. I am not looking for the most up to date features (although the more recent the better), I am just looking for reliable. Is there someone who produces a reasonably bug free distro and who keeps bug fixes updated reliably?
Dragonhead, from the LXF forums

A A distribution is purely a collection of numerous tools, utilities and services from discrete sources. It is practically impossible for anyone to produce a distribution which, by the time you install it, will not have any exploits or bugs. The best way to handle it is to install a recent distribution which will probably have the latest releases of core tools, such as RedHat 7.1, and then maintain the major programs, such as *Apache*, X, or whatever you will be using to suit your requirements.

All vendors, or at least the major ones, issue bug reports for almost everything on their distribution, at least when security is a factor, but it still relies on the third party source producing a fix for such problems in the first place. As most distros have some form of on-line package management tool, it is fairly straight forward to make sure that you have the latest version of a particular package.

The plugin architecture of *xmms* makes it flexible, but harder to configure.

Sound of silence

Q I have the following sound modules compiled (I had the same for kernel 2.2.18 and they worked fine):

sgalaxy.o, opl3.o, sound.o, soundcore.o, ad1848.o in the /lib/modules/2.4.2/kernel/drivers/sound directory. I made a symbolic link /lib/modules/2.4.2/misc to /lib/modules/2.4.2/kernel/drivers/sound so that sndconfig would work. I ran sndconfig as I did for 2.2.18 and both wav and MIDI were fine and I rebooted. However, when I use *xmms* to play a sound file, nothing comes from the speakers and it flies through the file much too fast. Any ideas?
Jimbo

»

«A Firstly, check to make sure that the modules are all loaded correctly using **lsmod**. If they are, try playing a mp3 with **mpg123 <filename>** and seeing what happens, as it may be a permissions problem with /dev/dsp or the use of a sound daemon, such as EsounD which is causing the problems. Assuming that works, check the configuration of *xmms* to ensure that it is using the correct plugin.

If you can't get any sound out of it at all, and the modules are not loaded, try to do **modprobe opl3** and then check /var/log/ messages to see why the module failed to load. The sound subsystem has not changed a great deal between 2.2 and 2.4, so the vast majority of modules should work in the same way. 'sndconfig' will set the 'sound-slot-1' alias in /etc/modules.conf, so by comparing the modules specified there with those which work should help solve any conflicts.

"We've been 'lovebugged' this week at work. Someone opened the virus on an oldish machine running Windows 98 and nuked the entire network"

Kernel query

Q I've just installed kernel 2.4, but on startup I get error message "mount fs type devpts not supported by kernel"

My question is, which configuration option allows support of devpts?

seaodyssey, from the LXF forums

A 'devpts' is a virtual filesystem which stores all the pseudo terminal devices, and is usually mounted on /dev/pts. You need to compile in support for 'Unix98 PTY support', which is usually kept at the default value of 256, as well as insert a line in to /etc/fstab to make sure that /dev/pts is mounted at boot time before anything can try to use the PTYs.

```
none  /dev/pts  devpts  gid=5,mode=620  0 0
```

A reboot of the machine into the new kernel should solve the problem.

No love lost

Q We've been 'lovebugged' this week at work. Someone opened the virus on an oldish machine running Windows 98 and nuked the entire network, sending lots of virus infected mails to some nice clients – oops!

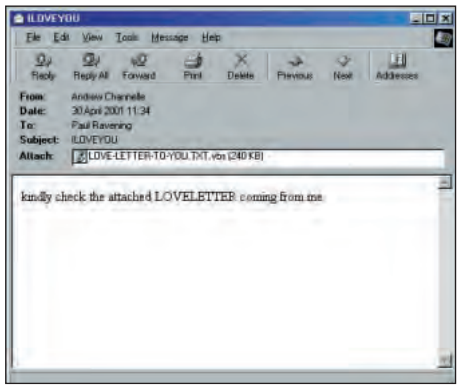

You should be able to stop the I Love You virus with something as simple as *procmail*.

We are using an old implementation of MS *Exchange Server* at the moment, and relying on most machines having their own up-to-date virus software (mostly McAfee/Norton). But as I said there are a few machines that are not protected, because the user has turned off the anti-virus software "cos it slows me pooter down." Most of the virus alerts are caused through the arrival of infected email.

Is there an effective way of detecting and dropping incoming email that contains any Script? And is it taxing on system resources?

Eric Howells

A Handling mail with Linux will let you send all mail through *procmail* which can easily be used to filter out nasty infections, spam and other unwanted email. A simple mail client, such as *qmail* or *postfix*, will handle the mail for a small to medium sized office on a low-end machine without any trouble. This does, of course, require an awareness of what you need to do in order to filter out specific viruses, as well as actually ensuring that the appropriate blocks are in place before it propagates to any of the internal machines. With something such as the 'LoveBug', it is quickest just to disable the script options within *Outlook*, which will stop any scripts being run automatically when the email is opened, rather than trying to struggle with anti-virus software and mail relays.

Selective KDE

Q I am running Mandrake 7.1 off of a *PCPlus* coverdisc. I want to upgrade *KDE* to version 2.1 from the *Linux Format* CD but don't know how! Obviously I don't want every language, nor do I necessarily want the development stuff, but what do I install, and in what order? Every file seems to be dependant on something else that I have not got yet! I tried *kpackage*, but maybe I should go with the command line, but the dependency problems will surely remain. Also If I replace files then, will I lose the old ones?

Mark Roberts

A With 'rpm', you don't need to install RPMs in any particular order if you try to install them as a batch. Simply doing: **rpm -Uvh /path/to/kde21/*rpm** as root, will return any dependency errors for individual packages before trying to install the whole set. If you have dependency problems, or are finding that Mandrake 7.1, which uses RPM version 3, doesn't like the RPMs on the CD, you may want to look at downloading compatible RPMs from www.rpmfind.net, or upgrading to Mandrake 7.2, which should allow you to use the RPMs on the CD.

KDE2 is certainly worth the trouble of upgrading, and it's not too hard to do either!

Get posting!

Join the *Linux Format* forum to get solutions to your Linux difficulties from knowledgeable users across the UK and beyond.
You can also let the LXF team know what they're doing wrong/right and what you'd like to see covered in the magazine and on the CD.

www.linuxformat.co.uk/forum.php

Don't panic!

Q I've just installed Mandrake 6.1 from a *PC Answers* CD. The install went fine, but now I can't get into Linux! I ran Rescue from the boot disk and it said "Unable to open an initial console – Kernel Panic: No init found. Try passing init= option to kernel." Can someone give some advice to a newbie like me?
Fishca, from the LXF forums.

A Aside from the fact that Mandrake 6.1 is a *very* old distribution which has been superseded by the 7.x, and even 8.x, releases, you are probably having problems with the rescue disk by either not passing the correct root partition in the 'root=' option, or by trying to boot an incorrectly installed distribution. You may want to try

```
linux root=/dev/hda5 init=/bin/sh
```

Of course, if you've got to the point of using the rescue disk, there is a more important problem which you've not explained, so reinstalling with a recent distribution, should solve any problems.

Directing traffic

Q Is there any way to set up BIND so that any.domain.com is directed to a certain port (ex.160) on domain.com? So instead of www.domain.com:160 it would be read www.any.domain.com.

A BIND, or rather named, only assigns hostnames to IPs, and not port numbers. HTTP, as default, runs on port 80 and this will be what the web browser will use to connect to the machine. You might want to look at using name based Virtual Hosting, that is, where a separate server configuration is used depending upon the hostname used to access the site, so 'www.any.domain.com' might serve a completely different site than 'www.domain.com'. If port 160 is to be forwarded to an internal web server, it is possible to use mod_rewrite to proxy requests to the external web server onto an internal machine with something similar to:

```
RewriteEngine On
RewriteRule /?(.*) http://10.1.1.2:160/$1 [P,L]
```

mod_rewrite will need to be compiled as part of *Apache* with the **—enable-module=rewrite**, or by building it as a DSO and having the existing Apache installation run it and the use of the **P** option to proxy the requests will require 'mod_proxy' which should be built in exactly the same way with **—enable-module=proxy**.

The big freeze

Q I've just compiled a new kernel from a tarball and the whole process seemed to go fine, but once I've copied the bzImage to the boot directory and used linuxconf to get lilo to point to the new kernel and then reboot the damn thing freezes at the first line.
Jeremy Inglis

A As LILO booted the kernel, then it is probably down to a mis-configured kernel. Start by getting the latest release, which is 2.4.4, or 2.2.19 if you're sticking to the 2.2 series, and start again from scratch with a kernel configuration. Generally, any configuration options which have been set when you did **make mrproper** should be left alone, unless you are 100% sure that you don't need them – such as SCSI options on a IDE only system. Once you have a machine that works, you can disable any options which you think you might need, one by one, then find out which one caused the machine to stop booting.

Fat pipe concerns

Q I'm getting an NTL cable modem connection next weekend. I need to know the best way to set up my network.

I have a RH7 server that acts as my Windows PDC, printer server and web server. I have Win95 and Linux (various flavours) clients.

I'd like to be able to expose my web server to the Internet in a secure fashion (i.e. using a firewall). I have an old 486 that I was going to install *SmoothWall* on, but from what I read I don't think that will allow me to expose the web server unless it is in a DMZ (third NIC). I don't have another spare machine to stick in a DMZ.

Is the best approach to use my existing RH7 server for all its current functions *plus* acting as the Cable Modem host? Or is there a way to use *SmoothWall* with my 486? If I go with the former, will this put the security of my network as risk, or can I tie things down properly without resorting to *SmoothWall?*
From Tenjin, on the LXF forums

A The most secure method is to get 2.4 running on the Red Hat machine, which should require a few updates, and use iptables to secure the machine and allow external services. You can then do something similar to;

```
iptables -t nat -A POSTROUTING -o eth0 -j
MASQUERADE
iptables -A INPUT -i eth0 —dport 80 -j ACCEPT
iptables -A INPUT -m state —state NEW,
INVALID -i
eth0 -j DROP
```

Which will allow connections on port 80, drop all other incoming data, and masquerade all internal machines behind the Red Hat box, assuming that the cable modem is on eth0. Any other services can easily be allowed or forwarded to an internal machine using the DNAT option with iptables. IPTables was discussed last issue, and you can find more information at http://netfilter.kernelnotes.org/

You will, of course, have to ensure that any services you run and allow external access to, such as *Apache*, mail and such like, are secure so that your gateway can not be exploited. It certainly isn't as neat having a gateway as its own DMZ, but it is better than not having any firewall at all.. LXF

Submission advice

We are happy to answer all sorts of Linux related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- **Please be sure to include any relevant details of your system. "I can't get X to work" doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.**
- **Be specific about your problem. Things like 'it doesn't work' or 'I get an error' aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?**
- **Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!**

We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

WRITE TO US AT:
Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email:
lxf.answers@futurenet.co.uk

Reviews

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

IN THIS ISSUE

PLUS

This month our panel of expert reviewers include **Richard Drummond, Paul Cavanagh, Tom Wilkinson** and **Jon Kent**

Top Stuff award

If we really, really like something, if it's the best in its field, then we'll give it our "Top Stuff" award. Only the very best will be chosen.

DISTRIBUTION

Progeny Debian 1.0

■ **LICENSE** Open source ■ **PUBLISHER** Progeny ■ **WEB** www.progeny.com

Debian is held to be the best – and purest – Linux distribution, but still has difficulty getting accepted by business. Richard Drummond discovers a distro that could change that.

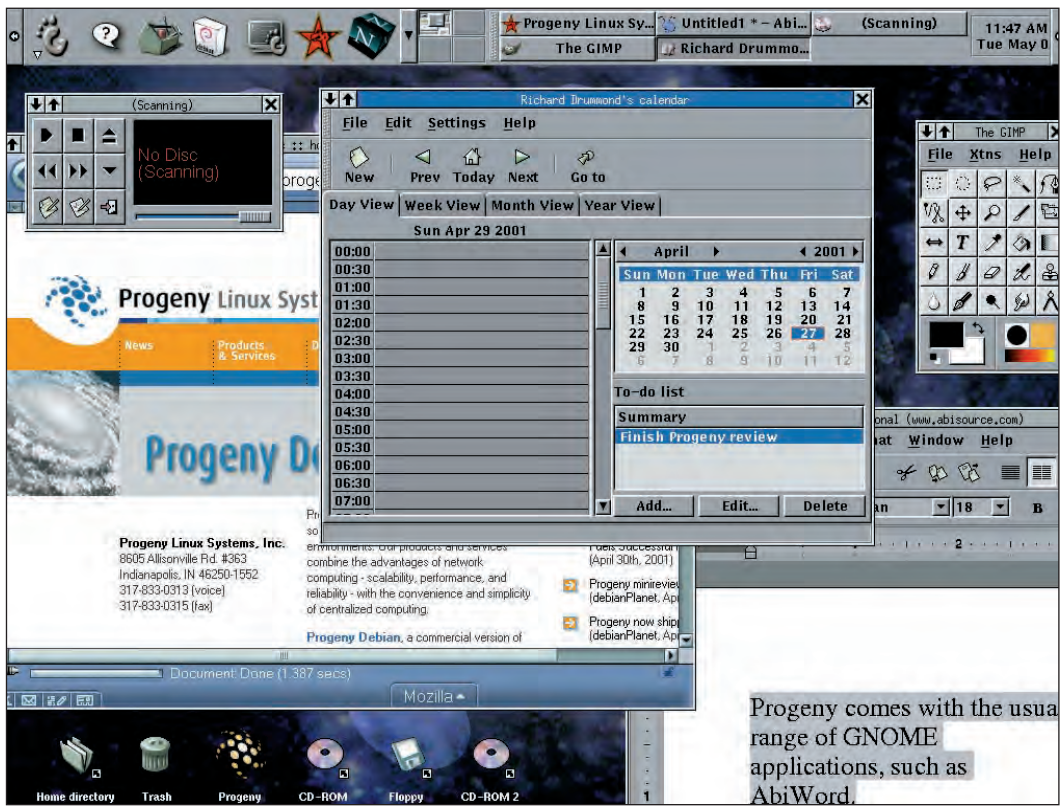

Two things are holding back Debian Linux's wide adoption in business: it's reputation for being hard to install and the lack of professional support. Progeny aim to solve both with this repackaging of Debian. What makes them think they can succeed where others have failed? Well, they are sticking close to Debian itself and they have Debian founder Ian Murdock at the helm.

Progeny Linux 1.0 is a two-CD distro based on Debian 2.2 with many, many software updates and some especially repackaged for Progeny. It has an entirely new installer, which performs hardware detection, and sports graphical configuration – but is still compatible with Debian 2.2.

Simpler installation

The Progeny installer presents you with a simple, graphical front-end to *GNU Parted* for partitioning. This is much easier to use than Debian's method, but is not without problems. *Parted* is fussy about partition tables and tends to take exception on tables that have been mucked about with by tools such as *Partition Magic*; and if it fails with an error, it causes the Progeny front-end to hang.

GNOME is Progeny's preferred desktop, but _KDE_ is also available.

Progeny offers you automatic partitioning – either using the entire disk or just the free space – or custom partitioning. The latter provides a GUI where you can edit your partition table, set mount points and so on. The non-destructive resizing of partitions is

More reviews in these sections...

Desktop 50
Jagged Alliance 2 – Among the endless first person shooters which pepper the Linux games market, can a turn based strategy game make any sort of impact? Paul Cavanagh finds out.

Professional 82
Solsoft NP Lite – Solsoft claims to make security and firewall management easy. We put it to the test.

Books – Another selection of Linux manuals.

Discovering hardware

Progeny's hardware detection scheme is based on Mandrake's detect library. Every time you boot, it probes for PCI and PCMCIA cards, USB devices, and SCSI and IDE drives. For each device it discovers, it tries to find a kernel module that will support it, and inserts that module into the kernel. The system knows how to do this because it has various predefined lists associating each device ID, where a driver is available, to a particular kernel module.

This scheme works quite well, but it does have a few peculiarities. No ISA PNP detection is done by default. That probably won't bother most people, but the problems with sound will. Some sound chipsets are mapped to OSS kernel modules and some to ALSA kernel modules. Progeny don't provide a build of ALSA for the kernels that they supply, however. The upshot is sound won't automatically be configured on some systems – such as, for example, the on-board sound in i810 and VIA686 chipsets. In these cases, you'll either have to build ALSA modules yourself or configure things manually. Another gripe is that the SCSI emulation layer is always used to drive ATAPI CD-ROMs – you don't get a choice.

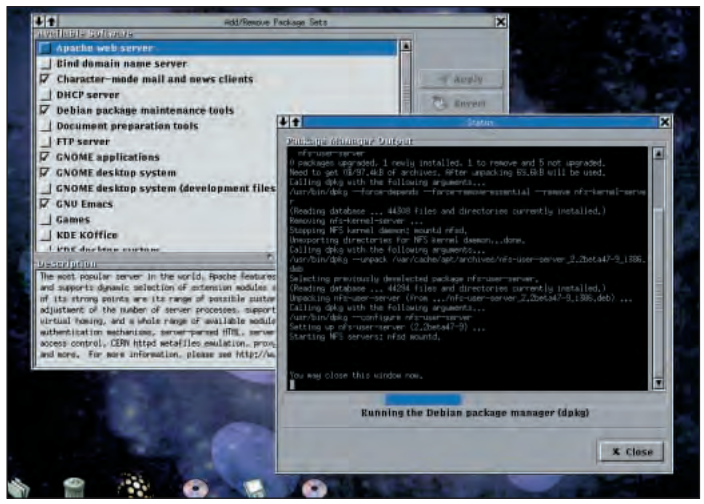

Progeny boast a new graphical front-end to package management.

not supported yet nor is ReiserFS. If you need to resize partitions, you can either switch to a console and use the normal text-based *Parted* interface, or do it the old-fashioned way with FIPS. Automatic partitioning just creates root and swap partitions: since Progeny uses GRUB as its bootloader, there's no need to have a boot partition.

Next, the installer copies the base system onto your drive and reboots into that and tries to configure an X server. If Progeny managed to detect your graphics card and mouse and your monitor is EDID-compliant, then everything is done automatically. Otherwise, you'll need to help out by choosing the appropriate options in the dialog provided. When that's done, Progeny installs X and *GNOME* and it's time for the real configuration to begin.

Debian provides a system called debconf which allows packages to display text-based configuration dialogs to the user when the package is installed or when reconfiguration is desired. As well as providing a *GNOME* front-end for debconf, Progeny supplements this with its configlets system – which uses Glade-designed, GTK-based dialogs to do basically the same job. During initial configuration, Progeny uses a set of six configlets: one each for time settings, user administration, mail delivery, network settings, printer settings and X configuration. (If you're already happy with your X set up, you can ignore the last one). These dialogs are also available as pages embedded in the GNOME Control Center on your desktop when installation is complete. Configuration is fairly self-explanatory, and tool-tips provide context-sensitive help should you get stuck.

Choosing software

Now you can choose the application software to be installed on your machine. Progeny make uses of Debian's excellent APT system for package management, but again presents a different user interface. In Debian, you may select packages for installation individually or by group as tasks. Progeny differs by grouping its packages into what it calls package sets, a more application-orientated way of look at things. Package sets offered include Apache, Emacs, Netscape, Mozilla, Samba and so on.

The program *gnome-apt-pkgset* lets you choose package sets to install and is the core GUI for package management on Progeny. However, there is no text-based alternative; there's no GUI-based way of choosing a package source to fetch packages from, and when you sync the list of available package against a remote list, you're are not told which packages have new versions available. It's much simpler and less bewildering for newbies than Debian's text-based *dselect*, but more coarsely-grained. Progeny runs gnome-apt-pkgset as the last phase of its installation, but it can be run at any time to install or remove software.

Apart from these graphical admin tools, living and working with Progeny is much the same as with Debian. It has a shiny, new theme for *Sawfish* and a special Progeny wallpaper, but it's still Debian underneath. Some of the default packages have changed, though. Progeny uses *Postfix* for mail delivery rather than exim, lprng rather than the plain BSD lpr for the printer daemon, and the easier-to use but less flexible *wvdial* dialup connections.

Debian for all?

Progeny Debian is an impressive first shot at a distro that will widen Debian's appeal. It doesn't stray too far the tried and tested formula, but the new features do make the system simpler to install, configure and manage – albeit, in some cases, at the price of flexibility. While there are some issues that need to be addressed, Progeny offers the power and unique flavour of Debian while still being accessible to newcomers to Linux. LXF

Support

By the time you read this, the boxed set of Progeny should be available to buy from the Progeny website. This includes a 400-page manual, 30 days of telephone support and a 90-day subscription to the Progeny Network.

The latter offers email support and notification of security updates, with more services planned for the future. Progeny also supply professional incident-based support for Progeny, Debian and other Debian-based distros.

Under the hood

Installer text/graphical
Package Format deb
Kernel version 2.2.18/2.4.2
Glibc version 2.2.1
Default compiler 2.95.2
XFree86 3.3.6/4.0.2
Default desktop GNOME 1.2

Linux Format VERDICT

Ease of use	8/10
Features	8/10
Performance	7/10
Stability	9/10

Takes the stability and power of the standard Debian Linux distro and boosts ease of use with simpler installation, graphical configuration and hardware detection.

LinuxFormatRating 8/10

DISTRIBUTION

Libranet Linux 1.9.0

■ **PUBLISHER** Libranet ■ **WEB** www.libranet.com ■ **PRICE** $20 (plus $5 shipping)

Have Libranet managed to simplify Debian for desktop users? Richard Drummond finds out.

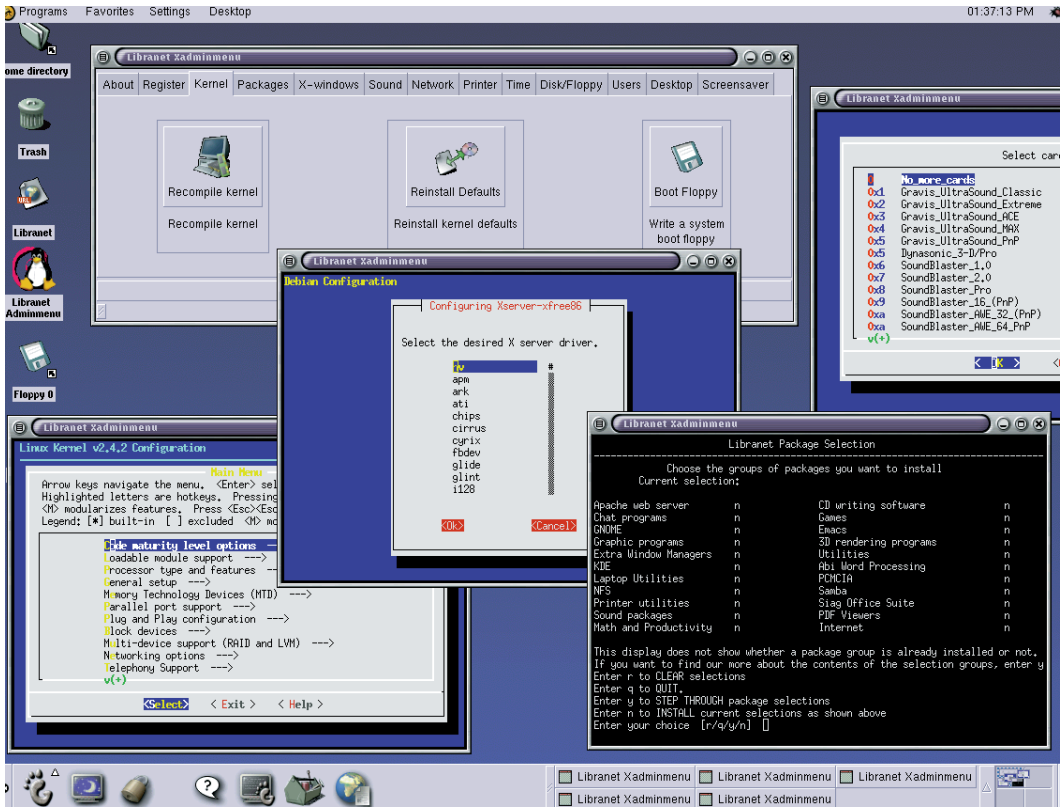

Libranet's innovative adminmenu speeds up day-to-day configuration tasks.

Debian is, of course, the best Linux distribution available but does pose some problems for the average desktop user. For example, its installer and packaging system are incredibly flexible but can be bewildering to the newcomer. Also, Debian's insistence on reliability leads to long delays between stable releases; if you want to keep up with the latest software versions you'll have to fetch packages from the testing or stable trees. Libranet attempt to solve these problems with their own repackaging of Debian.

This new version of Libranet, 1.9.0, is a single-CD version of Debian 2.2 (potato). While the latest stable Debian release (2.2r3) is powered by kernel 2.2.19 and XFree86 3.3.6 and has no *KDE* at all, Libranet's version boasts kernel 2.4.2, ReiserFS, XFree86 4.0.2 and *KDE 2.1*. It also includes Ximian's latest build of *GNOME*. If you need an up-to-date desktop, Libranet will save you a lot of downloading.

Debian simplified

Libranet uses the standard Debian two-phase install: install the base system, boot into that and then install and configure any extra software you need. The initial install on Libranet is done by a simplified version of Debian's menu-driven dbootstrap. The main difference here is that you can no longer insert kernel modules at this stage or set up your network. This means there's less to go wrong, but you can no longer install over a network. On the plus side, Libranet gives you the option of installing to a ReiserFS partition – but you'll still need to create an Ext2 boot partition, since LILO cannot boot from a Reiser partition. And like Debian, you'll either have to make your partitions beforehand or use the spartan cfdisk.

The second install phase is entirely different in Libranet. It presents you with a list of software available for you to install, and you must step through the list saying to whether you want them installed. It's crude but it works, and once you get everything installed you can always fall back on using Debian's dselect. It then installs the packages you want, and begins configuration of your network card, sound card (with ALSA), modem and X server. Like Debian 2.2, Libranet does no hardware detection. With the exception of the network settings, everything is configured with the standard debconf dialogs. In essence it's not easier than standard Debian, but at least you are led through the steps in order and forced to get everything sorted before you begin.

Aside from the newer software, the Libranet desktop offers no really big changes over the standard Debian. The one novel feature is the Adminmenu, a single control centre from which you can launch a host of configuration and admin jobs. Here you can configure and rebuild the kernel; install new packages; reconfigure your X server, network settings, soundcard or printer; mount and format floppies; tweak your desktop and more. In the main, everything is done though the usual console-based config dialogs that Debian provides. Package installation is done from the Libranet CD with the same primitive installer as before, or you can download packages from the web by entering their names. The latter, though, is a completely useless option – I don't see why they didn't provide some menu-driven means of setting up an apt source (like Debian's apt-setup) and then use a standard apt front-end like dselect or gnome-apt for selecting packages.

In the balance

Libranet is less ambitious and not so slick as some of the other reworkings of Debian, but is still a useful distro despite (or perhaps because of) that. By not venturing too far from the proven Debian formula, Libranet have created a solid distro with a good selection of desktop software. You can, if you want, use it like a vanilla Debian install, using all the usual admin and package tools. However, for less able users, the Adminmenu makes things easier – and is a great time-saver for everybody. Having said that, the Libranet package selector desperately needs to be replaced. There are some other missed opportunities, too. For instance, the newer builds of XFree 4.0.2 from the Debian woody and testing trees support the automatic detection of video cards, input device and monitor frequencies; and the latest kernel packaging system supports GRUB as a boot loader. Including these in Libranet would make it a much easier to use distribution. LXF

Linux Format VERDICT

Ease-of Use	7/10
Features	8/10
Performance	8/10
Value for money	9/10

A cheap and cheerful remix of Debian Linux targeted firmly at desktop users.

LinuxFormatRating 7/10

IMAGE MANIPULATION

Photogenics 5.0

■ **PUBLISHER** Paul Nolan ■ **WEB** www.paulnolan.com ■ **PRICE** $99

There is a new challenger to The GIMP throne. Tom Wilkinson takes a look at a new multi-platform graphics package that could unleash the artist in every Linux user.

As the Linux operating system grows in popularity, more software manufacturers are waking up to the new markets which are starting to present themselves. Among these are Paul Nolan, author of *Photogenics*, which is being released to three platforms: Linux (Intel), Windows and AmigaOS. The program, originally developed for AmigaOS, has been ported to Linux and windows using an in house cross-platform toolkit, similar to wxWindows. There are however disadvantages, the biggest of which is speed. On our test machine, a 650MHz PC with 128MB RAM, some of the dialogue boxes, most notably the "file open" box, were very slow and flickery while expanding the directory tree. This could have something to do with the fact that it also generates previews of *all* the files before opening at the same time, which is no doubt a pretty intensive operation.

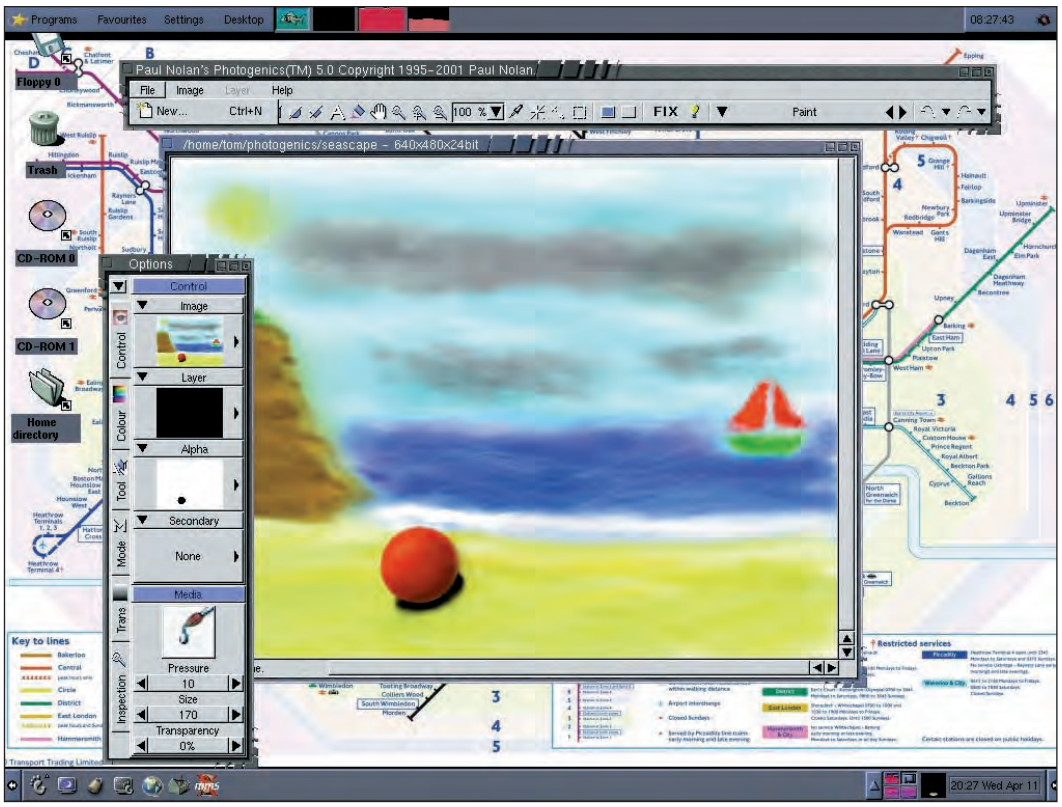

The Photogenics interface. The icons are located in one of two floating windows in the desktop.

Installation

Installing the program is relatively straightforward – a shell script is provided which will check your system is able to run the software, and then start the installation procedure proper, by asking for the directories to install the binaries and data files. It's worth noting that the installation program has to be run as root unless you only want to install in your home directory. This aside, however, the installation was straightforward, with the exception of one thing. The installation instructions suggest the program will be added to the menus of either *GNOME* or *KDE*, should they be installed. However, neither of these additions were apparent on our system due to a quirk of Mandrake, which uses a separate menu system across all desktop environments rather than leaving it to each window manager.

On loading, the interface is presented as something of a cross between *The Gimp* and *PhotoPaint*. While the toolbox and the graphics are presented in their own floating windows, there is still a windows-style icon bar across the top, along with menus for the most common operations, such as resizing and file operations. This is a good compromise between having everything in a single window as is common with most Windows packages, and having lots of toolboxes all over the desktop.

The manual is supplied in HTML format. However, there is very little formatting information present, and this leads to *Netscape* rendering the page in a very small font which can be extremely hard to read. The manual fares little better in *Konqueror* or *Mozilla* – the former has problems with some of the graphics included while the latter has similarly tiny text. I can't help feeling that the pages have been devised for *Internet Explorer* and only hastily adapted for Linux. There are also references to Windows-only features. Fortunately it is well written and easy to follow.

Features

The opening page for the manual in *Photogenics* states that it has an "incredibly powerful and flexible method of working". This is certainly true. The program introduces a number of new concepts which help a great deal in the production and modification of images. Perhaps the most significant of these is the way that the undo function works. Unlike other paint packages, where two mouse buttons are bound to painting colours, only a single colour can be selected at any time in *Photogenics*. While at first this may sound like a step backwards, the function the right button now performs is far superior to simply painting, and is, in fact, a logical extension of what many people were using the button for in the first place: undoing. This is highly useful because it allows the underlying pattern of the image to be restored far more easily and without having to resort to a traditional undo option – something which can be fairly infuriating in other paint packages due to the amount of work which may be undone with a single click. In addition, there is an icon which will clear all the modifications you've made since you last selected the paint colour, or brush size. This helps if you've made a complete mess of things. It's worth noting, though that the right mouse button does not produce a simple eraser effect, but performs the exact opposite of the left mouse button, so can be used to lighten textures or to tweak an area that has been painted in ways that would simply be impossible in other packages.

While the concept of Layers is not new to graphics programs, the way in which *Photogenics* implements and

uses them is somewhat different. Upon selecting a new brush type, size, pressure or colour, a new layer is always created. This helps to facilitate the undo feature mentioned above – the remainder of the image is kept, safe and sound, separately from the changes on the base layer. Along with this comes the alpha channels (see Boxout), which can be used to determine which area of the image can be painted on, and how much this paint will show through. These are innovative, in that while most other graphics packages support only monochrome alpha channels, *Photogenics* supports full colour ones, adding an extra degree of flexibility.

The program also offers a far wider range of drawing media than most of the competition. These include painting modes that are designed to simulate real art materials such as watercolours, pencils, and airbrushes. The results are effective and, with a little judicious varying of the pressure setting, can produce some very realistic results indeed.

On top of the natural media features, there is a whole set of "pyrotechnic" drawing tools including three different types of fire, lens flare, and neon. Each of these tools can be used in conjunction with one of the many paint modes to produce some innovative and interesting effects. Each of the masterpieces gracing these two pages were produced, at least in part, using this highly useful feature.

The brush off

With "Experimentation Mode" switched on, selection of a new brush, paint colour or special effect will replace the one you have just been using. This is a departure from the conventional method of generating special effects – the area to which the effect is applied is selected using an ordinary paintbrush tool rather than by drawing a shape around the area to be affected. Naturally, the undo function of the right mouse button is still available here. It's also worth noting that if your PC is not fast enough to draw these effects directly, the experimentation mode will allow you to draw the area that you wish to affect with a different painting mode, for example green paint, and then replace these areas with the special effect you want to use. To finalise a modification while staying in experiment mode, you press the "Fix" button on the toolbar.

This brings me on to one of my major gripes with the interface – the Fix button is placed right next to the undo button on the toolbar. More than once, when attempting to correct a mistake, I pressed Fix and was lumbered with unwanted changes. Similarly I have pressed undo when meaning to Fix. While *Photogenics* may have a very flexible method of undoing mistakes, on a very fundamental level there is still only one level of undo. Once you've fixed a modification, you're stuck with it. In addition, should you move a layer, then fix your modifications, the new layer is created exactly where you left the old one – this means that there are a number of areas you won't be able to paint to until you move the layer back, which can be annoying. Regular saves are a vital tool when using this package.

The most annoying part of the program, though, was the intermittent problem I experienced when zooming – the program would often crash completely if I attempted to change the magnification level. Coupled with the lack of any sort of autosave feature, this could get extremely infuriating. I was, however, unable to estimate when the program would fall over and when it would continue working properly. It's entirely possible that this is a problem unique to my setup – but I doubt it.

The colour selection also felt like it occasionally had a mind of its own; giving a different shade to the one selected. The problem was most noticeable when trying to select white and consistently getting a light grey. This was not a palette problem, as far as I could tell.

In and out

Photogenics supports over 20 image file formats, but supports both loading and saving of only half that number. This might be a problem for a small number of users, but is not too much of a worry. The majority of the rest are save-only filters for exporting to formats suitable for desktop publishing such as encapsulated PostScript. The usual trio of JPEG, GIF and PNG are all available to read/write, and these are probably what most users will need for images to be published on the web. LXF

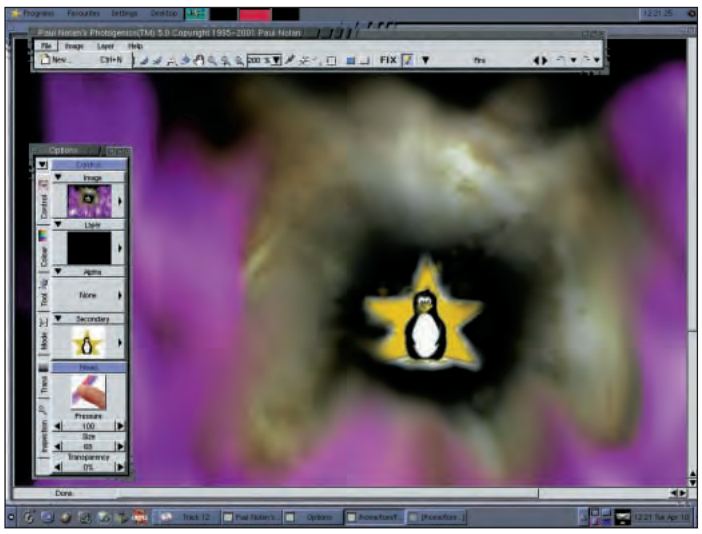

Various painting effects are available, including a rub-though to a separate image.

Alpha Channels

One of the most important parts of any modern graphics package is it ability to create and use alpha channels. This can be thought of as an extension to the simple stencil which allows paint to cover some areas but not others. Alpha channelling goes one step further with this, and allows a defined amount of colour through the stencil depending on the value of the filter at that point (0 is opaque, 255 totally transparent). Photogenics goes beyond this and allows separate red, green and blue alpha channels to be used, for example only letting certain colours through in specific areas.

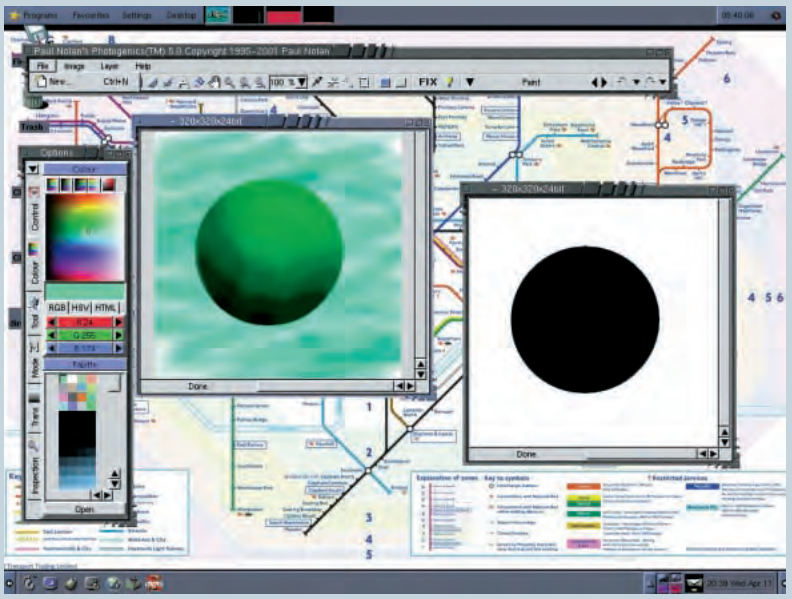

Alpha channelling in action. The background was painted in afterward using the alpha channel (left) to block the paint from the ball.

Linux Format VERDICT

Ease of use	8/10
Features	9/10
Performance	7/10
Value for money	8/10

***Photogenics* is a well rounded graphics package that shows great promise. Only a couple of problems prevent it getting a perfect score.**

LinuxFormatRating 8/10

Roundup

Every month we compare tons of software, so you don't have to!

Our selection at a glance

- Fltdj
- Korganizer
- slpim
- Jpilot
- Plan
- GNOME-PIM

PIMs

Nick Veitch dumps his filofax and assesses whether Linux really can sort out your life for you.

The computerisation of the workplace means that, for better or for worse, contact books and rollodex are more and more a thing of the past (I don't know about you, but with all the computers I have, there isn't even room on my desk for a Rollodex). And seeing as a lot of business communication takes place through the computer these days, is it any wonder that many people now have some form of computerised organiser.

Managing contacts and appointments using your Linux box should theoretically be easier than using a filofax – you are unlikely to forget where you left your computer. A good organiser will combine some sort of calendar function, a task list, and a contact management system to store numbers, email addresses or indeed anything else you might need to know. It should also be able to cooperate with your email client, and perhaps with a PDA too, for those moments when you are away from your desk.

All the PIMs covered here are available on the CD apart from *Korganizer*, which is part of the general KDE distribution. If you have *KDE2* installed, you probably already have *Korganizer!*

Surprisingly, there isn't quite as much choice when it comes to Linux PIMs as you might expect. There are the usual *GNOME* and *KDE* contenders, but very few others, and most lack the common features and functionality you might expect from a fully fledged PIM. Many of the ones featured in our roundup have taken different approaches to the task, and it would be hard to truthfully describe any of them as being particularly similar to any others. But the real question is, can any of them sort your life out for you?

The PIMs in this roundup have been rated on their ease of use, features and interoperability with other common desktop software. It may be that some of these criteria are not so important to your situation, so please, read the individual reviews before making up your mind!

Fltdj

■ **VERSION** 0.4.2 ■ **HOMEPAGE** www.geocities.com/letapk/linux.html

The name comes from The Daily Journal, and the fact that this application makes use of the FLTK – Fast, Light, Tool Kit for its interface. The interface also employs the FLEditor extension to create editable text dialogues, used for the appointments and ToDo list.

As is so often the case, this project is unfinished, and there are quite a few limitations that would prevent anyone but the most undemanding people finding it competent or comprehensive enough for daily use.

You can store contact details in *fltdj*, but once again these are in the most simple format, consisting of just one text string – there is no actual database of telephone numbers or email addresses, and no way of sharing or gathering this information from other applications.

One nice feature is that you can use the program as a diary as well – each day can have a text entry of up to 1000 characters, which is enough for the "went to work, worked, went home to bed" style of journal keeping.

There is a limited amount of support for alerts. All appointments are automatically alerted (but only if fltdj is running – there is no alert daemon process to warn you of upcoming events) and you will be alerted for each entry, assuming you have entered a valid time. This is one area that could do with some work, though there are a few nice features. Particularly useful is the ability to adjust the alerts to suit the appointment – there's no point being warned that you are supposed to be in a meeting hundreds of miles away at the time the meeting takes place!

The advantages of fltdj are that it is comparatively easy on resources (runs in just over 1Mb), and has a bright and friendly style. The calendar doesn't take up too much screen space (though it would be nice to have a few options here, like turning the clock off), and since everything is plain all entries are searchable. Though there is no import/export facility, the file format used is so simple you could easily knock up a script to add or extract useful data.

Linux Format VERDICT

Ease of use	9/10
Features	4/10
Integration	1/10

A lightweight organiser for those who don't have a hectic lifestyle.

LinuxFormat**Rating** 4/10

Korganizer

■ **VERSION** 2.1 ■ **HOMEPAGE** http://korganizer.kde.org/

KDE has made great progress in providing an integrated and easy to use desktop environment. You might expect the 'official' *KDE* organiser to be up there with the best in terms of features and ease of use.

It is also unsurprising that *Korganizer* borrows heavily from other popular organisers for Windows/Mac and other Desktop OSes. When you run the software, you'll be presented with a three-way split view – a main window showing today's appointments, a smaller view with this month's calendar and a summary of your to do list. You can then change the view to view by week, get a detailed list of to dos and so on.

Adding an event or appointment is fairly straightforward, and the software supports recurring events for birthdays, monthly reports and so on. Your appointments can be grouped into different categories, which can be displayed in different colours if you want to set these up in the configuration dialog.

Each day is blocked out into half hour slots, but you can have events overlapping, in which case they are displayed side by side. There are quite a few clicks required to set up a meeting, but you don't need to fill in all the data if it's not relevant, and fortunately you can reschedule meetings, and adjust duration, simply by dragging and resizing the box in your daily/weekly display.

The *Korganizer* daemon will add itself to the kicker, to make sure that you receive any alerts you have set, and also providing a convenient place to launch the application from.

Documentation is a problem though. Although there is enough online help to get you started, the details of the software certainly aren't explored in depth. *KDE* applications tend to be well supported with on-line help, but only up to a point, and this is one of the cases where the documentation isn't sufficient to do the software justice.

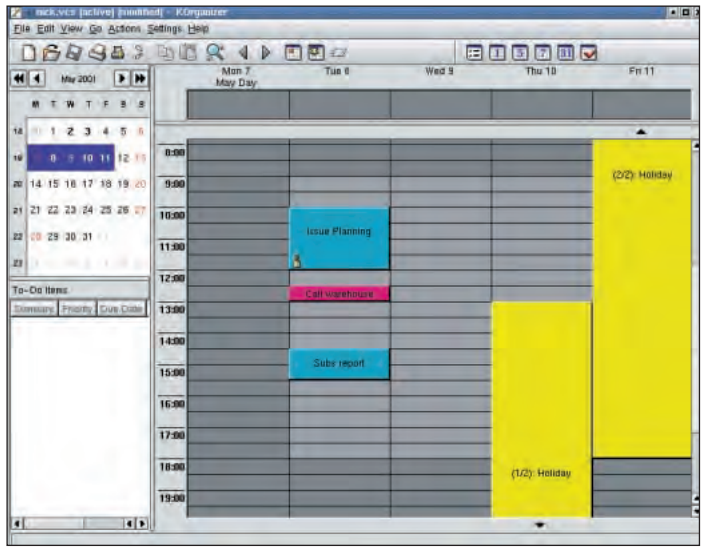

Appointments can be moved around the views by dragging with the mouse. Specific colours can be set up for types of event.

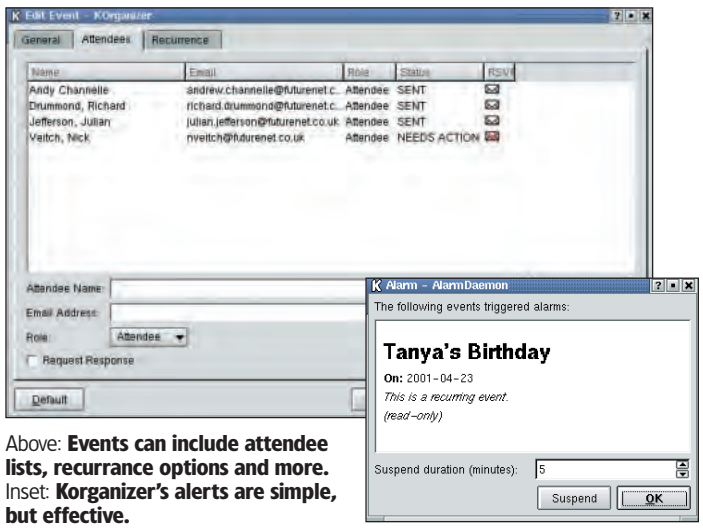

Above: **Events can include attendee lists, recurrance options and more.** Inset: **Korganizer's alerts are simple, but effective.**

However, the degree of localisation is impressive. As you might expect, *Korganizer* is available in around 40 languages. What's more impressive is that it has the national holidays for over twenty countries programmed in. It might have been nicer if you could display some of these simultaneously, but it's still very useful. All these days are automatically counted as holidays as far as the organiser is concerned, so they'll be greyed out on your 'working week' view.

Support for other file formats is limited, but you can import calendars directly from *Ical*. Contact management is actually handled via Kab, another component of the *kdepim* package. This attempts to provide a uniform structure for contact information, though it isn't directly supported by many other applications, even the current version of *kmail*, though this is likely to be fixed. Kab information includes all the standard vCard information and can convert/import data from Eudora, Netscape and MS Messenger formats with the help of *Kmailcvt*.

Kab can be launched directly through the Address Book button when adding attendees to the list. Sadly the integration with *kmail* to send out emails isn't working fully yet, but this will be a handy feature, hopefully by the next big update.

Korganizer is certainly feature rich compared to the other PIMs on test here, but could do with better integration with other packages, better documentation, and a few tweaks to make the array of features work well for the user.

Linux Format VERDICT

Ease of use	8/10
Features	8/10
Integration	4/10

Good features, but it needs better integration and more documentation.

LinuxFormatRating **7/10**

slpim

■ **VERSION** 0.0.6 ■ **HOMEPAGE** http://slpim.sourceforge.net/

If there is a black sheep of the PIMs family, this is it. Can you imagine using a terminal-based PIM? Well, the author certainly can and, to be fair, there's no reason why it shouldn't work, you'll just have to learn a *lot* of keyboard shortcuts.

Unfortunately, development on this project hasn't got very far yet, and as such it is merely an address book rather than a full-blown PIM. Further modules are planned, and who knows, it could turn out to be another really useful terminal-based app. There is certainly a case for this approach – you only have to consider how, for example, *ncftp* can be much more effective for downloading some specific files from your regular ftp haunt than using a graphical client like gftp or *Konqueror*.

As it stands it can quite successfully be used as an address book for *mutt*, and additional modules, particularly a calendar, will definitely enhance its usefulness.

It is, as you might expect, very fast, but if you are looking for useability, features or sanity, look elsewhere for the moment.

Linux Format VERDICT

Ease of use	4/10
Features	1/10
Integration	2/10

Works as a *mutt* phonebook, but is not much use otherwise.

LinuxFormatRating **2/10**

Jpilot

■ **VERSION** 0.99
■ **HOMEPAGE** www.jpilot.org

Jpilot's main function, as you might infer from the name, is to co-ordinate your diary, tasks and address book with your Palm Pilot device. But as well as simply downloading info from your Pilot, *Jpilot* can be used to edit the info, make new appointments, add entries to the todo list and more – in fact, you don't even really need a Palm Pilot to make good use of it.

If you do have a Palm Pilot, the latest version will attempt to connect to the device to sync your data. The software supports USB devices as well (providing you have an up-to-date kernel), and merely requires that you press the HotSync button on the device to get it connected.

Obviously, the features in the software are limited to the common features of the standard Pilot software, though. You shouldn't find this too much of a restriction.

Enter an appointment by clicking on a day, then adjusting the spin gadgets to the relevant times. The events can then have various notes attached and can be set to repeat on specific days, weekly, monthly or yearly. There are no definable categories for different types of event, but you can have events with no particular times attached (useful for birthdays etc) and these can have alarms set to remind you of where you should be at any given time. The alarms, by the way, only work when the data is in your Pilot, not through the Linux desktop.

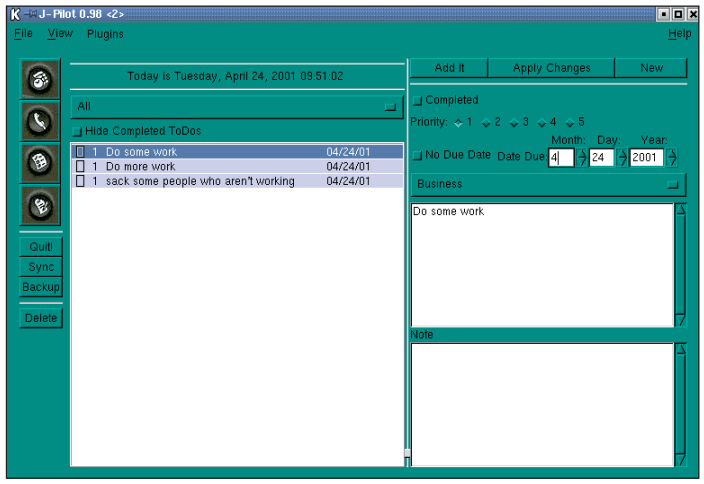

Connectivity is JPilot's strong point. All your data can be simply synchronised with a Palm Pilot.

The contact manager is pretty good, and allows you to store name, address, several contact numbers/email addresses, and includes custom fields for adding additional information. You can also set up different categories here for separating the business and personal details of people in your system.

To do lists are also supported with facilities to add items, cross them off when they are done, and even set time limits for their completion. The list view can be used to hide items which have already been done – handy if your list, like mine, usually comes to several dozen items a day.

Jpilot's functionality can also be extended with the use of plugins, of which there are some useful examples. *Jpilot Mail* is a plugin which enables you to download pending mail from the Pilot and send it via your desktop, and also store mail received via the desktop to the Pilot.

While obviously it makes sense to have a Pilot PDA to use this, *Jpilot* is a pretty good choice for a basic, low-overhead organiser.

Linux Format VERDICT

Ease of use	7/10
Features	6/10
Integration	4/10

Works supremely well in conjunction with the Palm Pilot.

LinuxFormatRating 6/10

Aethra

■ **VERSION** beta 2
■ **HOMEPAGE**
www.theKompany.com

Billed as unified messaging and groupware, *Aethra* is based on *Magellan*, and is still in a fairly raw state. Although intended to support Calendars etc in the future, these features are not yet complete.

The contact manager is pretty good though. It contains all the data you would expect from a vCard compatible system, but with some nice touches for easy entry. A good example here is the pop-up menu for international dialling codes when entering a phone number. Select the correct country and the +xxx prefix is added to the field.

This is a very flash looking project, but is currently devoid of almost any useful features you might expect, bar contact management. It does occupy a serious amount of memory though, being even more bloated than *Korganizer*.

Linux Format VERDICT

Ease of use	6/10
Features	2/10
Integration	2/10

Bloated and incomplete, but holds promise for the future.

LinuxFormatRating 2/10

Plan

■ **VERSION** 1.8.4 ■ **HOMEPAGE**
www.bitrot.de/plan.html

Plan has been around for a really long time. Development started back in '93, so it's no surprise that this program is now pretty stable, easy to use, and full of nice touches.

In terms of functionality, *Plan* only claims to be a calendar application. There are no to do lists, contact management or other bits bolted on here, but apart from these, it is the most complete in terms of calendars.

One of the best features is the multiplicity of views available. Although these do not dynamically resize, there are two sizes of monthly view, a week, day, year and year overview display. The latter is like a linear year planning strip, which you can scroll backwards and forwards through.

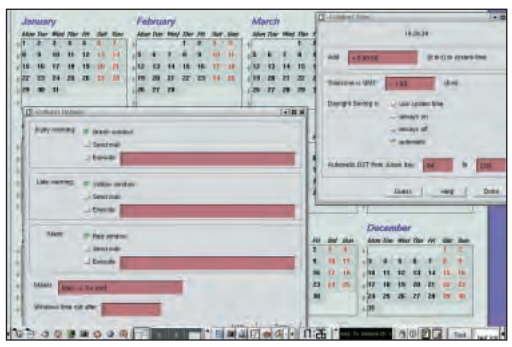

A venerable contender with unique calendar function.

Events you have entered are displayed in coloured blocks, and you can specify the usual repeating properties, and even set up shell scripts to be executed. *Plan* has an alert daemon to remind you of things you need to be reminded of, which will pop up a suitably scarey window at the appropriate time.

You can configure your own holidays by simply entering the information in a text editor. While this may seem a little like too much hard work, *Plan* has an intelligent parser for the file which understands dates like 'easter' and terms like 'first Thursday in October'.

Plan also works with a network server, to co-ordinate calendars for group use. Item lockouts mean that no two people can edit the same appointment simultaneously, so data integrity should always be maintained.

Some aspects of the user interface are a little fiddly, but you can, for example, resize and move events by dragging the boxes in any of the views.

While it misses out on many features, it is comprehensive in terms of event planning, and incredibly light on system resources.

Linux Format VERDICT

Ease of use	8/10
Features	5/10
Integration	1/10

Small and fast, but missing a few needed features.

LinuxFormatRating 5/10

GNOME-PIM

■ **VERSION** 1.4.0 ■ **HOMEPAGE** www.gnome.org/gnome-office/gnome-pim.shtml

GNOME PIM is actually a small collection of packages, notably *GnomeCal* and *GnomeCard*. *GnomeCal* opens up with the usual triple view of monthly calendar day view and to do list. Appointments and To Do tasks can be added, the view can be changed to weeks, months and years, and all the other options you might expect.

Individual appointments can be rescheduled by dragging with the mouse, but only on the day view, which is a bit disappointing – to move a meeting to another day requires you to edit the details, but at least the pop-up calendar display makes it easy to change the dates.

GnomeCal has about the most comprehensive settings for alerts available in our roundup. You have the choice of displaying a message, playing an alert sound, running another application or sending out an email – all at separately definable times. Even some of our staff would be hard pressed to forget something with that many reminders. Sadly though, this feature highlights one of *GnomeCal's* shortcomings: the lack of integration with other apps. To send reminder emails you have to manually type in an email address – why not sync up with *GnomeCard?*

The To Do list includes the ability to prioritise entries, and various categories of tasks. These have due dates and associated notes too – the list will automatically highlight overdue items if you like, and an unnerving countdown of remaining time would have even some LXF writers handing in their work on time.

Sadly, *GnomeCal* suffers from the same lack of documentation as *Korganizer*. There is some help, but it really only skims through the basics.

GnomeCard is pretty comprehensive when it comes to keeping tabs on people's information. There is much more implemented here than is defined in the basic vCard format, including stuff which varies from superfluous (latitude) to handy (PGP keys). *GnomeCard* has minimal documentation though, and although it's handy for storing info, it doesn't link to any applications which might be able to make use of it.

As with *Korganizer*, this package is let down only by its lack of integration and documentation. Feature-wise it could be a little better: linking to *GnomeCard* for a contact's email details would be handy, as would linking to your email client.

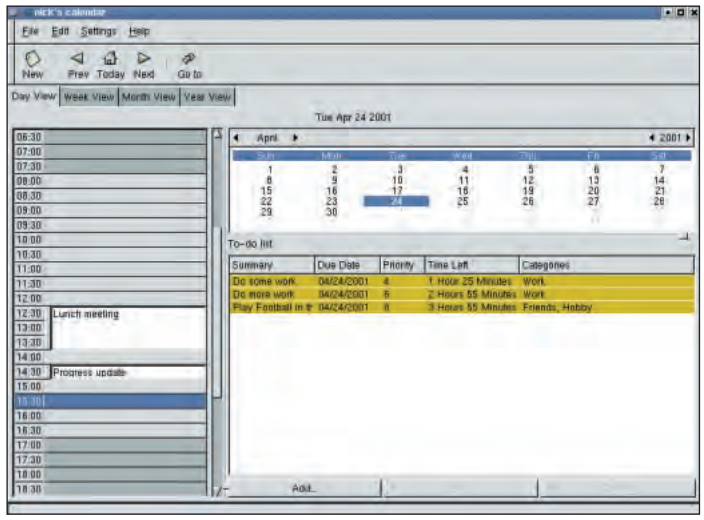

The basic view shows your daily schedule, a calendar and to do list.

Linux Format VERDICT

Ease of use 7/10
Features 7/10
Integration 1/10

Could do with more docs, more features and better integration.

LinuxFormatRating 6/10

PIMs THE VERDICT

Although not as high profile as many of the other office applications, there are a variety of PIMS and PIM-like applications around for Linux, and certainly enough variety to be sure of offering something for everyone.

The two foremost do-it-all applications we have unearthed are *Korganizer* and *GNOME-PIM*. That doesn't mean you should write the others off though. In terms of being easy to use, compact and economical on system space, *Jpilot* has much to recommend it, especially if you have a Palm Pilot.

If your contact management and mail facilities are currently handled adequately elsewhere, you may feel that Plan's lack of these features is not so much of a problem, while its ability to share data across a network through the simple use of shared files is a facility you could make use of in virtually any networked environment.

All of the programs were lacking integration features to some extent. *Korganizer* relies on kab for example, to create a unified resource for contact management – a good thing. Unfortunately *Korganizer* is about the only thing which currently uses it.

For a general PIM solution, *Korganizer* and the *GNOME-PIM* programs are probably your best bet. In terms of features, *Korganizer* may be slightly ahead of the *GNOME* offering, at the expense perhaps of being slightly more complicated to use in places – sometimes having too many options and settings just makes simple tasks take longer. *GnomeCal's* alerts and todo lists are better thought out, but the calendar itself is harder to use and less visually rich.

Both could do with some much better documentation, and both are pretty heavy on m resources, so if you want a solution for your laptop, *Jpilot* or *Plan* may be a better option.

Table of features

	FLtdj	Korganizer	slpim	Jpilot	Plan	Aethra	GNOME-PIM
Journal entries	y	n	n	y	n	notes	n
To Do lists	y	y	n	y	n	notes	y
Contact Management	limited	kab	y	y	n	y	y
Calendar	y	y	n	y	y	n	y
Public Holidays	n	y	n	n	y	n	n
Reminders	y	y	n	†1	y	n	y
Types of event	n	y	n	n	n	n	y
Import formats	n	vcal, ical †2	n	†2	n	†2	vcal
Palm Sync	n	†3	n	y	n	n	n
Data merge	n	y	n	†1	y	n	n
Toolkit	fltk	kde/qt	slang	GTK+	Motif	kde/qt	Gnome
Memory usage *	1.2Mb	9Mb	600k	4Mb	2.5Mb	11Mb	8Mb†4
Overall score	4/10	7/10	2/10	6/10	5/10	2/10	6/10

* NB it is impossible to calculate exact memory usage, as this depends on a number of factors, not least of which being the configuration of the machine you are running on. These figures are a guide only.

†1 Only on Palm device †2 Other conversions through additional utilities

†3 Not yet implemented fully †4Combined for GnomeCard and GnomeCal

HotPicks

The best new open source software on the planet!

This is the place where we get to profile some of the hottest software around.

Each month we trawl through the hundreds of open source projects which are released or updated, and select the newest, most inventive and best for your perusal. Most of the Hot Picks are available on our CD, but we've provided web links if you want to make sure you have the very latest version.

If you have any suggestions for things that we should cover, email us at linuxformat@futurenet.co.uk.

HotPicks at a glance

Hot Picks awards

Everything covered in our Hot Picks section is unmissable, but every month we'll be singling out one project for outstanding brilliance. Only the very best will be chosen!

AL SIMULATOR

Achilles

■ **VERSION** 0.0.4 ■ **WEB** http://achilles.sourceforge.net

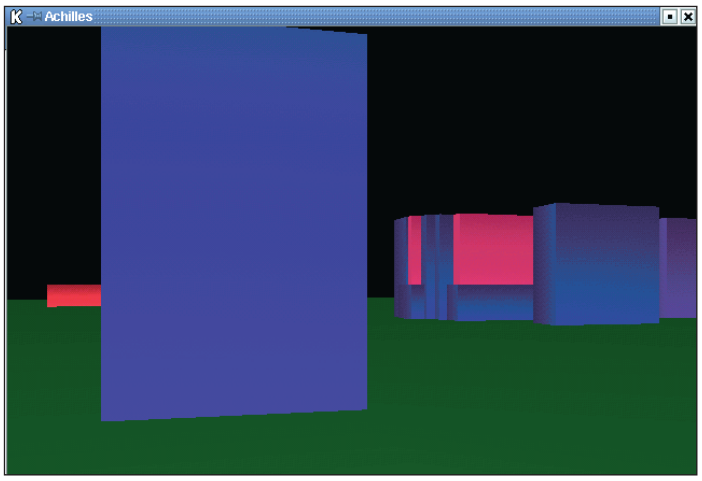

Achilles beings inhabit a distinctly cubist universe.

Artificial Life is the creation and study of complex systems which exhibit some of the properties of organic life. One of the favourite avenues of research in AL is the use of computers to create virtual universes populated with artificial organisms that behave like real ones – they eat, fight, breed and evolve. In such system, the wetware of real life is substituted for software. A classic AL is experiment along these lines is Larry Yaeger's *PolyWorld*. *Achilles* is a less sophisticated multi-platform version of the PolyWorld universe.

As in *PolyWorld*, an organism's behaviour in *Achilles* is controlled by its "brain", a Hebbian neural network. The composition of that neural network and various physical characteristics of the organism are determined by its genetic code. *Achilles*' organisms are simpler, however. In Polyworld, organisms are polygonal; in Achilles they are cuboid. In Polyworld they have a distinct front and can see in the direction they are looking. *Achilles*' organism don't need to learn how to see; they use some kind of telepathy to detect food and their fellow creatures.

Simpler they may be, but *Achilles* organisms do exhibit life-like behaviour. An organism may move around, eat food when it is hungry and interact with the other denizens of the universe. Eating food will replenish its energy and restore its health. Taking damage during a fight will decrease its health. When its health hits zero, an organism dies and becomes food for the rest. The inhabitants here have no difficulties with cannibalism.

When two creatures meet their neural nets determine whether they will fight or mate. If its the latter, then the genetic code of any offspring is a fusion of the parents' DNA with some random mutations thrown in for luck.

The *Achilles* universe is visualised in 3D using OpenGL. The organisms are cuboid blocks of varying sizes and colours that scuttle around on the flat surface of the world, while food is represented by yellow blocks. Interactions between creatures are represented by laser-like flashes. A blue flash signifies mating, a red one fighting and a yellow one eating.

Your role in this limbic-centred world is that of a passive observer. You can move the position and height of your camera's-eye view on the universe and pause the action.

The *Achilles* binary accepts various command-line options to set the initial conditions of the universe: its size, the initial number of organisms and the amount of food available. One annoying thing is that the settings for the initial number of creatures and food are also minimum settings; if the number of either in the universe falls below the minimum then an organism or some food will be spontaneously generated and positioned at random. While this means that life will never die out, it spoils the natural dynamics of the simulation.

A feature that is completely missing from *Achilles* is some kind of statistics on the population and evolution of the universe. How long has it existed? What's the current number of organisms? What's the average life-span? And so on. The code is fairly simple, however, so it would be fairly easy for somebody to add this sort of thing.

Achilles may seem entirely pointless to you; if that's so, then simply regard it as virtual fishtank and watch it running for stress relief. If, however, you have any interest in AL, it's an intriguing piece of software.

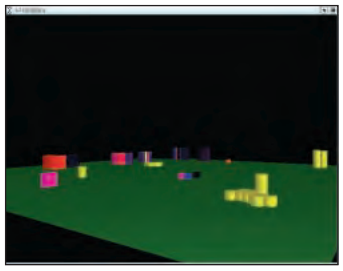

Your role in the whole affair is that of an observer.

While the *Achilles* universe is less complex than *PolyWorld* and its denizens less obviously life-like, it clearly demonstrates one of the principles of AL: that complex behaviour can emerge bottom-up in a rule-based system.

In a NUTSHELL

Ideal if you have an interest in artificial life, but otherwise, it's just eye candy.

PERSONAL FINANCE MANAGER

Gnucash

■ **VERSION** 1.4.12 ■ **WEB** http://www.gnome.org/

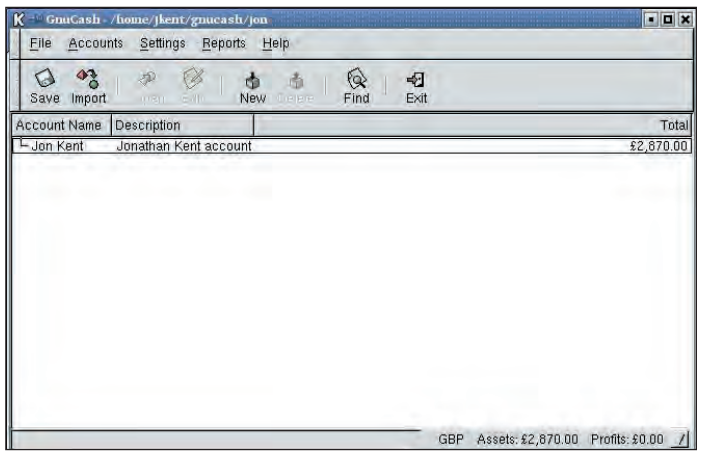

This is the opening screen, where you can choose an account to view.

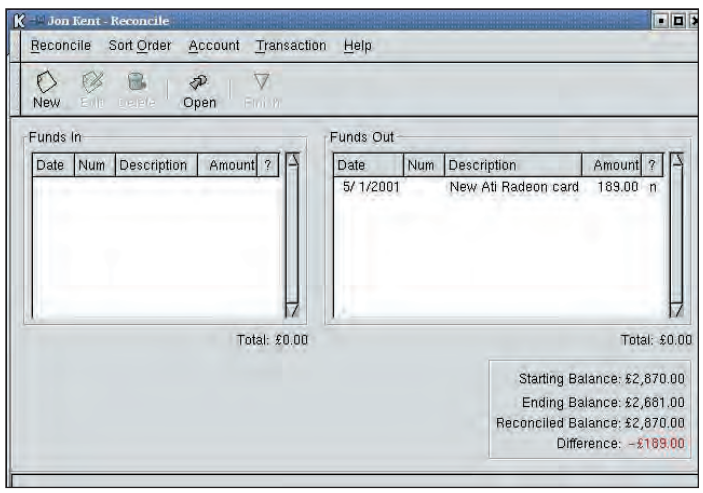

You can see which transactions are imminent in the reconcile view.

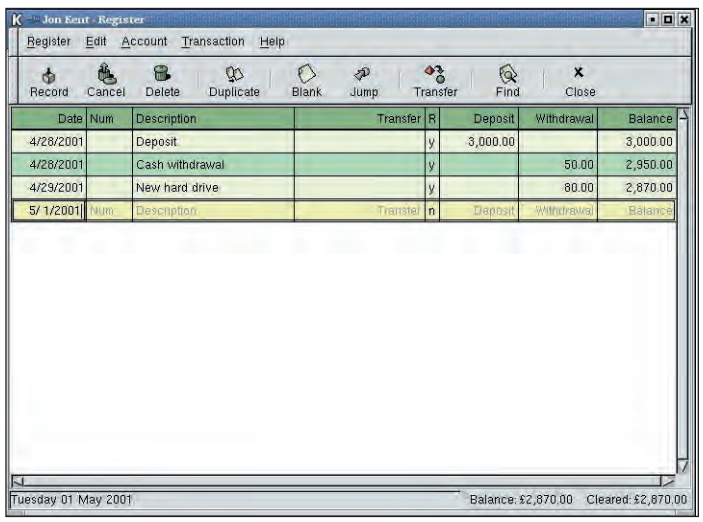

The register's look and feel will be familiar to Quicken users.

Over the past months here at HotPicks we have shown you several applications that competently replace their Windows counterparts. In many cases, these applications are far better than their commercial rivals. The only significant area that we have not yet covered is personal finance, so to rectify the situation here is *Gnucash*, an open source personal finance manager. *Gnucash* forms part of the *GNOME Office* suite, but like all other *GNOME Office* applications works perfectly well under other window environments.

Gnucash allows you to manage all of your finances, as long as you are prepared to put time into inputting the details of all transactions and setting up new accounts as required (it can't do everything!). The interface is clean and logical and getting starting is simple enough. If you have previously used *Quicken*, you can even import your old Quicken files.

To define a new account you simply click on the New icon and a window pops up asking you to enter a few details, such as the type and name of the account and a description. Once you have set up your account, these details will appear on the main window. To access the account you simply double-click on the one you wish to open and this will display the Register window – which will be instantly familiar to anyone who's used either of the two big Windows personal finance packages.

The Register, or Ledger, window is where all transactions, such as deposits and withdrawals are logged. To enter a new transaction on the account you simply enter the date of the transaction, a brief description and the total amount in the appropriate column. One column of interest is the R, or reconciled, column which can be set to yes or no. This allows you to note that a cheque has been issued against the account, but that it has not be drawn. Once the cheque is cashed, you would change the reconciliation status to yes.

As *Gnucash* allows you to have both reconciled and unreconciled items logged against your account, the Register window displays at the bottom of the screen two balances: Balance and Cleared. Balanced is the state of the account assuming all items are reconciled and Cleared shows the state of the account without unreconciled items. This feature is excellent for planning a weekly budget, as it provides a snapshot of the account as it is, and as it will be when all payments have cleared.

To reconcile an account you select Accounts and Reconcile from the menu and enter a statement date, or accept the default, which is the current date. A new window will be displayed that shows transactions that need to be reconciled. You do this by selection the transaction you want to reconcile and double clicking to change the status to yes. At the bottom of the window the various balances and any differences between them are displayed. As you would not reconcile transactions until you get you bank statement you should be able to reconcile all transactions and not end with a balance difference. Should some disparity occur – and the laws of the universe almost guarantee it – you would then need to manually check why there is a difference.

In addition to these functions *Gnucash* has several standard reports than can be run against the accounts setup such as Profit and Loss, Balance Sheet and so on. You can also write your own reports, but you will need to know Scheme in order to do so.

As you may have noticed *Gnucash* requires a lot of manual intervention, and there does not appear to be a method by which you can enter standard monthly comings and goings from your accounts such as direct debits or salary payments. This means that every month all transactions have to be entered. A tedious process at the best of times. It is also unfortunate that you have to know Scheme in order to generate a custom report, this should really be configurable within the GUI. However, *Gnucash* is a very stable product and worth a look at if you need a personal finance package.

In a NUTSHELL

A basic, but useful personal finance package, that is very easy to use thanks to its good help documentation. It is a shame that custom reports require programming and there is no method to enter ongoing transactions.

DOCUMENTATION EDITOR

manedit

■ **VERSION** 0.4 ■ **WEB** http://wolfpack.twu.net/ManEdit/

Man pages are the backbone of any UNIX system, and everyone expects to be able to run man against any application that they are working with. However, there are a lot of man pages missing or out-of-date and this needs to be addressed urgently. No one pretends that writing documentation is easy, and up until now they were a complete pain. But here comes *manedit* to the rescue. It does exactly what the names implies, allows easy creation and updating of man pages.

When you run *manedit* and open up a page, you notice that the text is not displayed with the groff syntax. This is because manedit converts each page section from groff in Manual Page XML format and, when the page is saved, converts the file back into groff and saves it. This conversion process was chosen as the developers felt that groff is too overly complex to write directly in, and therefore chose to use the Manual Page XML format as a middleman to page creation.

Back to the actual work of creating a man page. Creating a new page is simply a case of pressing the New icon and selecting whether to use a template or use a completely blank page and add you own sections and headings. The templates provided should meet most requirements, with versions for api, config, intro and program man pages. Once you select the template you wish to use, you are presented with two panes within the manedit window. The left pane displays the section headings and the right pane will display the contents of the current section heading. If you wish you can delete and add sections as you require and even cut and paste details from other man pages. A nice little feature is the ability to preview your page to ensure the layout and details are correct.

The help available with the program, which itself just references man pages, is well written and will get you up and running in next to no time. As mentioned, even though the editor uses a XML format, there are very few tags that you need to learn, and if you use a template you could just write the details in the tags already laid out.

manedit is a must-have application. If you ever wanted to give something back to the Linux community, but felt that you could only do this by programming, consider installing this application and volunteer your time to update or create man pages for scantily documented projects, it is of equal importance as code writing.

In a NUTSHELL

Very good man page editor with a unique approach to formatting which negates the need to learn groff. Stable and well thought out, a vital app.

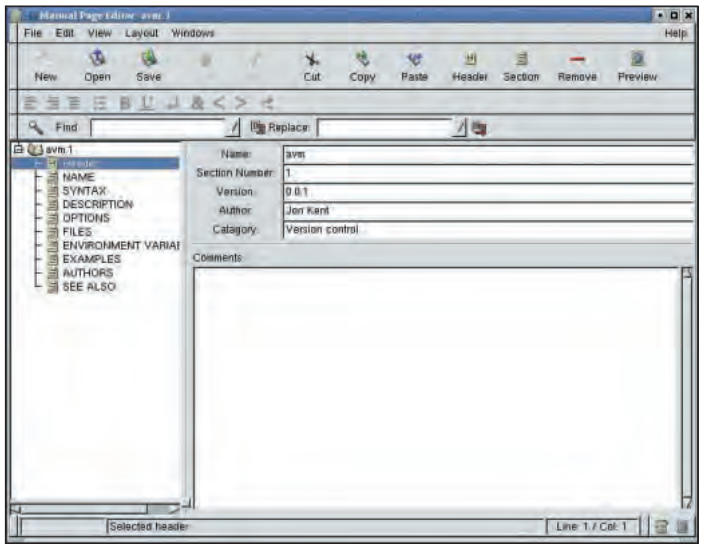

manedit simplifies the creation of man pages, and no need for groff. This is the main view where most of the work is done.

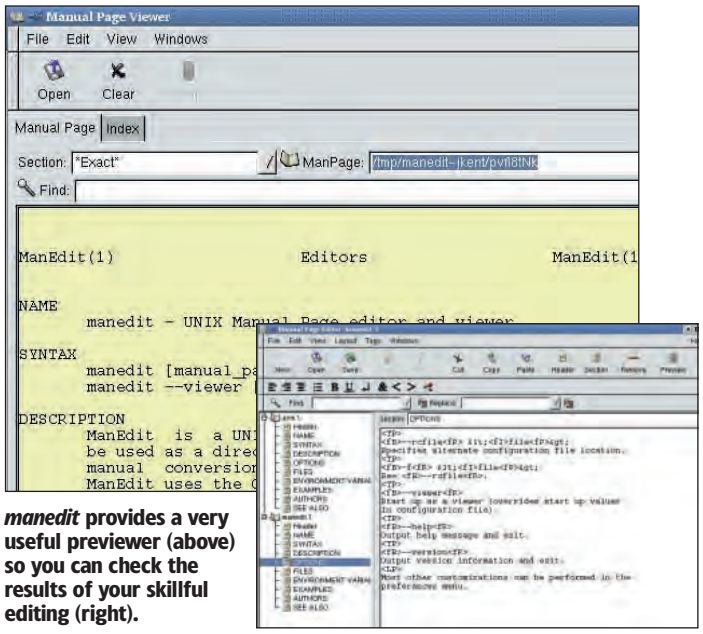

***manedit* provides a very useful previewer (above) so you can check the results of your skillful editing (right).**

NETWORK CLIPBOARD

Netclip

■ **VERSION** 1.0
■ **WEB** http://netclip.freax.eu.org

Picture the scene. A colleague at work asks you for the phone number or email address of a contact, which you duly paste into an email and send to him or her. Simple. But could it be simpler? After all, this colleague is on the same network as you. Wouldn't it be much quicker if there were some way you could more directly copy and paste between yours and your colleague's desktops. Well,

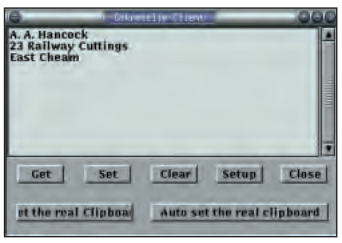

It's like copying and pasting, but across the network.

there is. This is what *Netclip* allows – it implements a shared clipboard which can be accessed over a network.

Netclip is, naturally enough, implemented as a client-server pair. If you just need one shared clipboard, you just need to run one server on your network. Everybody who needs to access it must run a netclip client. Clients exist for GTK+, Windows, and a *Kylix* client is also available. The package is currently available only as source code.

Netclip provides minimal documentation, but its trivially easy to use. Each client need only be told the IP address of the server maintaining the shared clipboard. Each client provides a simple text entry gadget where you can enter and view text. Click 'Set' and the contents will be posted to the shared clipboard; click 'Get' and the contents of the clipboard will be retrieved and displayed. Not surprisingly, you can paste from the local clipboard to the text gadget as per normal, and the interface provides a switch to automatically update the local clipboard with the contents of the client's text entry area.

Installation can be a bit more tricky, though. If you want the Netclip server launched automatically during boot-up, you'll have to copy and link its start-up script into your system's startup-scripts.

As with any network service, *Netclip* is a security risk – especially since it offers no means of authenticating users, but, the author claims, that should be addressed in a future release.

In a NUTSHELL

***Netclip* is an incredibly simple idea, but is an effective way of speeding up collaboration within the network.**

DEFENDER CLONE

Defendguin

■ **VERSION** 0.0.6
■ **WEB** http://www.newbreedsoftware.com/defendguin

As you may have guessed from the name of this little game, *Defendguin* is a *Defender* clone with a nice little twist. The aliens are in fact none other than a bunch of little Bill Gates clones who are trying to capture and mutate the Tux population of the planet. Your mission is to ensure that all the Bill Gates get terminated, and that you rescue as many Tuxs as possible. Now this is FUN!

There is a great deal of satisfaction seeing one of the Bill aliens blasted into a million pieces.

The program use the Simple DirectMedia Layer which is designed to be a cross-platform library that provides access to the underlying framebuffer and audio devices. If you wish to install this game you will need to pop over to the web site for the software at http://www.libsdl.org/. The links to this site are also available on *Defendguin's* web site.

As with the original *Defender* game you have a laser and a limited amount of bombs that can be used to blast all the alien Bills currently viewable on the screen. Aside from the fun of blasting up Bill Gates-a-likes, the game itself is great fun to play and a quite addictive. The graphics, whilst not exactly of *Quake III* quality, do the job, and the whole games runs at a very respectable speed. For a game that the author states is still in alpha, it is very stable and no major problems where found our test system.

In a NUTSHELL

A good *Defender* clone that allows you to shoot Bill Gates aliens to bits, what more could you want? Good speed and okay graphics make this a very addictive game that will keep you busy into the early hours.

***Defendguin* is a clever *Defender* clone that casts you in the role of a weather-beaten penguin protector.**

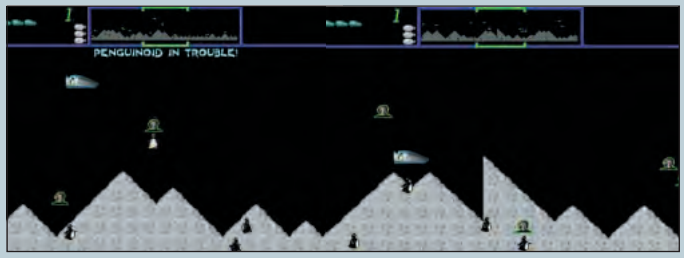

DOCUMENT PROCESSOR

TeXmacs

■ **VERSION** 3.3.15 ■ **WEB** http://texmacs.org

Donald Knuth, one of the pioneers of Computing Science, is perhaps best known for the creation of *TeX*, the mathematical type-setting package (see What on Earth, issue 1). Now, *TeX* is a great piece of software, but, even with a macro package such as *LaTeX*, it's not the easiest thing to get to grips with.

TeXmacs is much more accessible, providing a WYSIWYG interface to *TeX*, much more like a modern graphical word-processor. The second-half of the name is obviously derived from *Emacs* but don't let that worry you. While you can program in Guile Scheme, *GNOME's* Lisp-like language, *TeXmacs's* core functions are all available via point and click.

TeXmacs looks and feels a lot like any word processor. You can enter text, select text with mouse, cut and paste to the clipboard, choose font faces, sizes and styles and so on. You can also insert images (vector formats like EPS and XFig are handled) and hyperlinks to documents. In conjunction with *ISpell* it even supports spell-checking. But it does a whole lot more.

Like *LaTeX, TeXmacs* comes with built-in document styles – such as letter, article and seminar – to ensure your documents have a consistent and professional look. Within a document, you can choose style packages – such as basic, theorem and program – to typeset certain extracts of your document. Each document style defines things like how paragraphs are formatted and how sections of a document and marked, and you can write your own styles. *TeXmacs* styles are not compatible with *LaTeX* styles, but you can import and export *LaTeX* format documents.

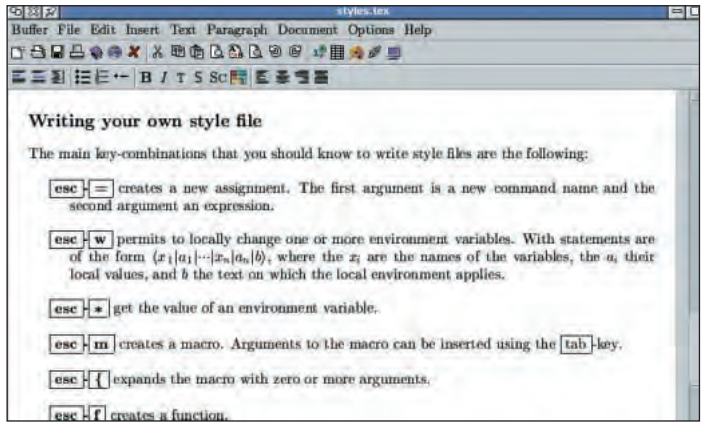

The real power of *TeXmacs* lies in its tools for inserting formulae and equations and for creating and editing tables. When entering formulae, for example, pressing ALT and a Latin character, gives its Greek equivalent. Very useful. Then there's a range of mathematical symbols and functions for inserting, numbering and formatting common structures like definitions, lemmas and theorems. Add to that the features for creating and managing bibliographies, tables of contents and lists of figures, and you have a package which is much more adept at handling technical documents than your average word-processor. LXF

***TeXmacs* is an excellent word processor for creating technical or mathematical documentation.**

In a NUTSHELL

A great word processor for technical and mathematical documents, that is quicker and simpler to use than wither *TeX* or *LaTeX*.

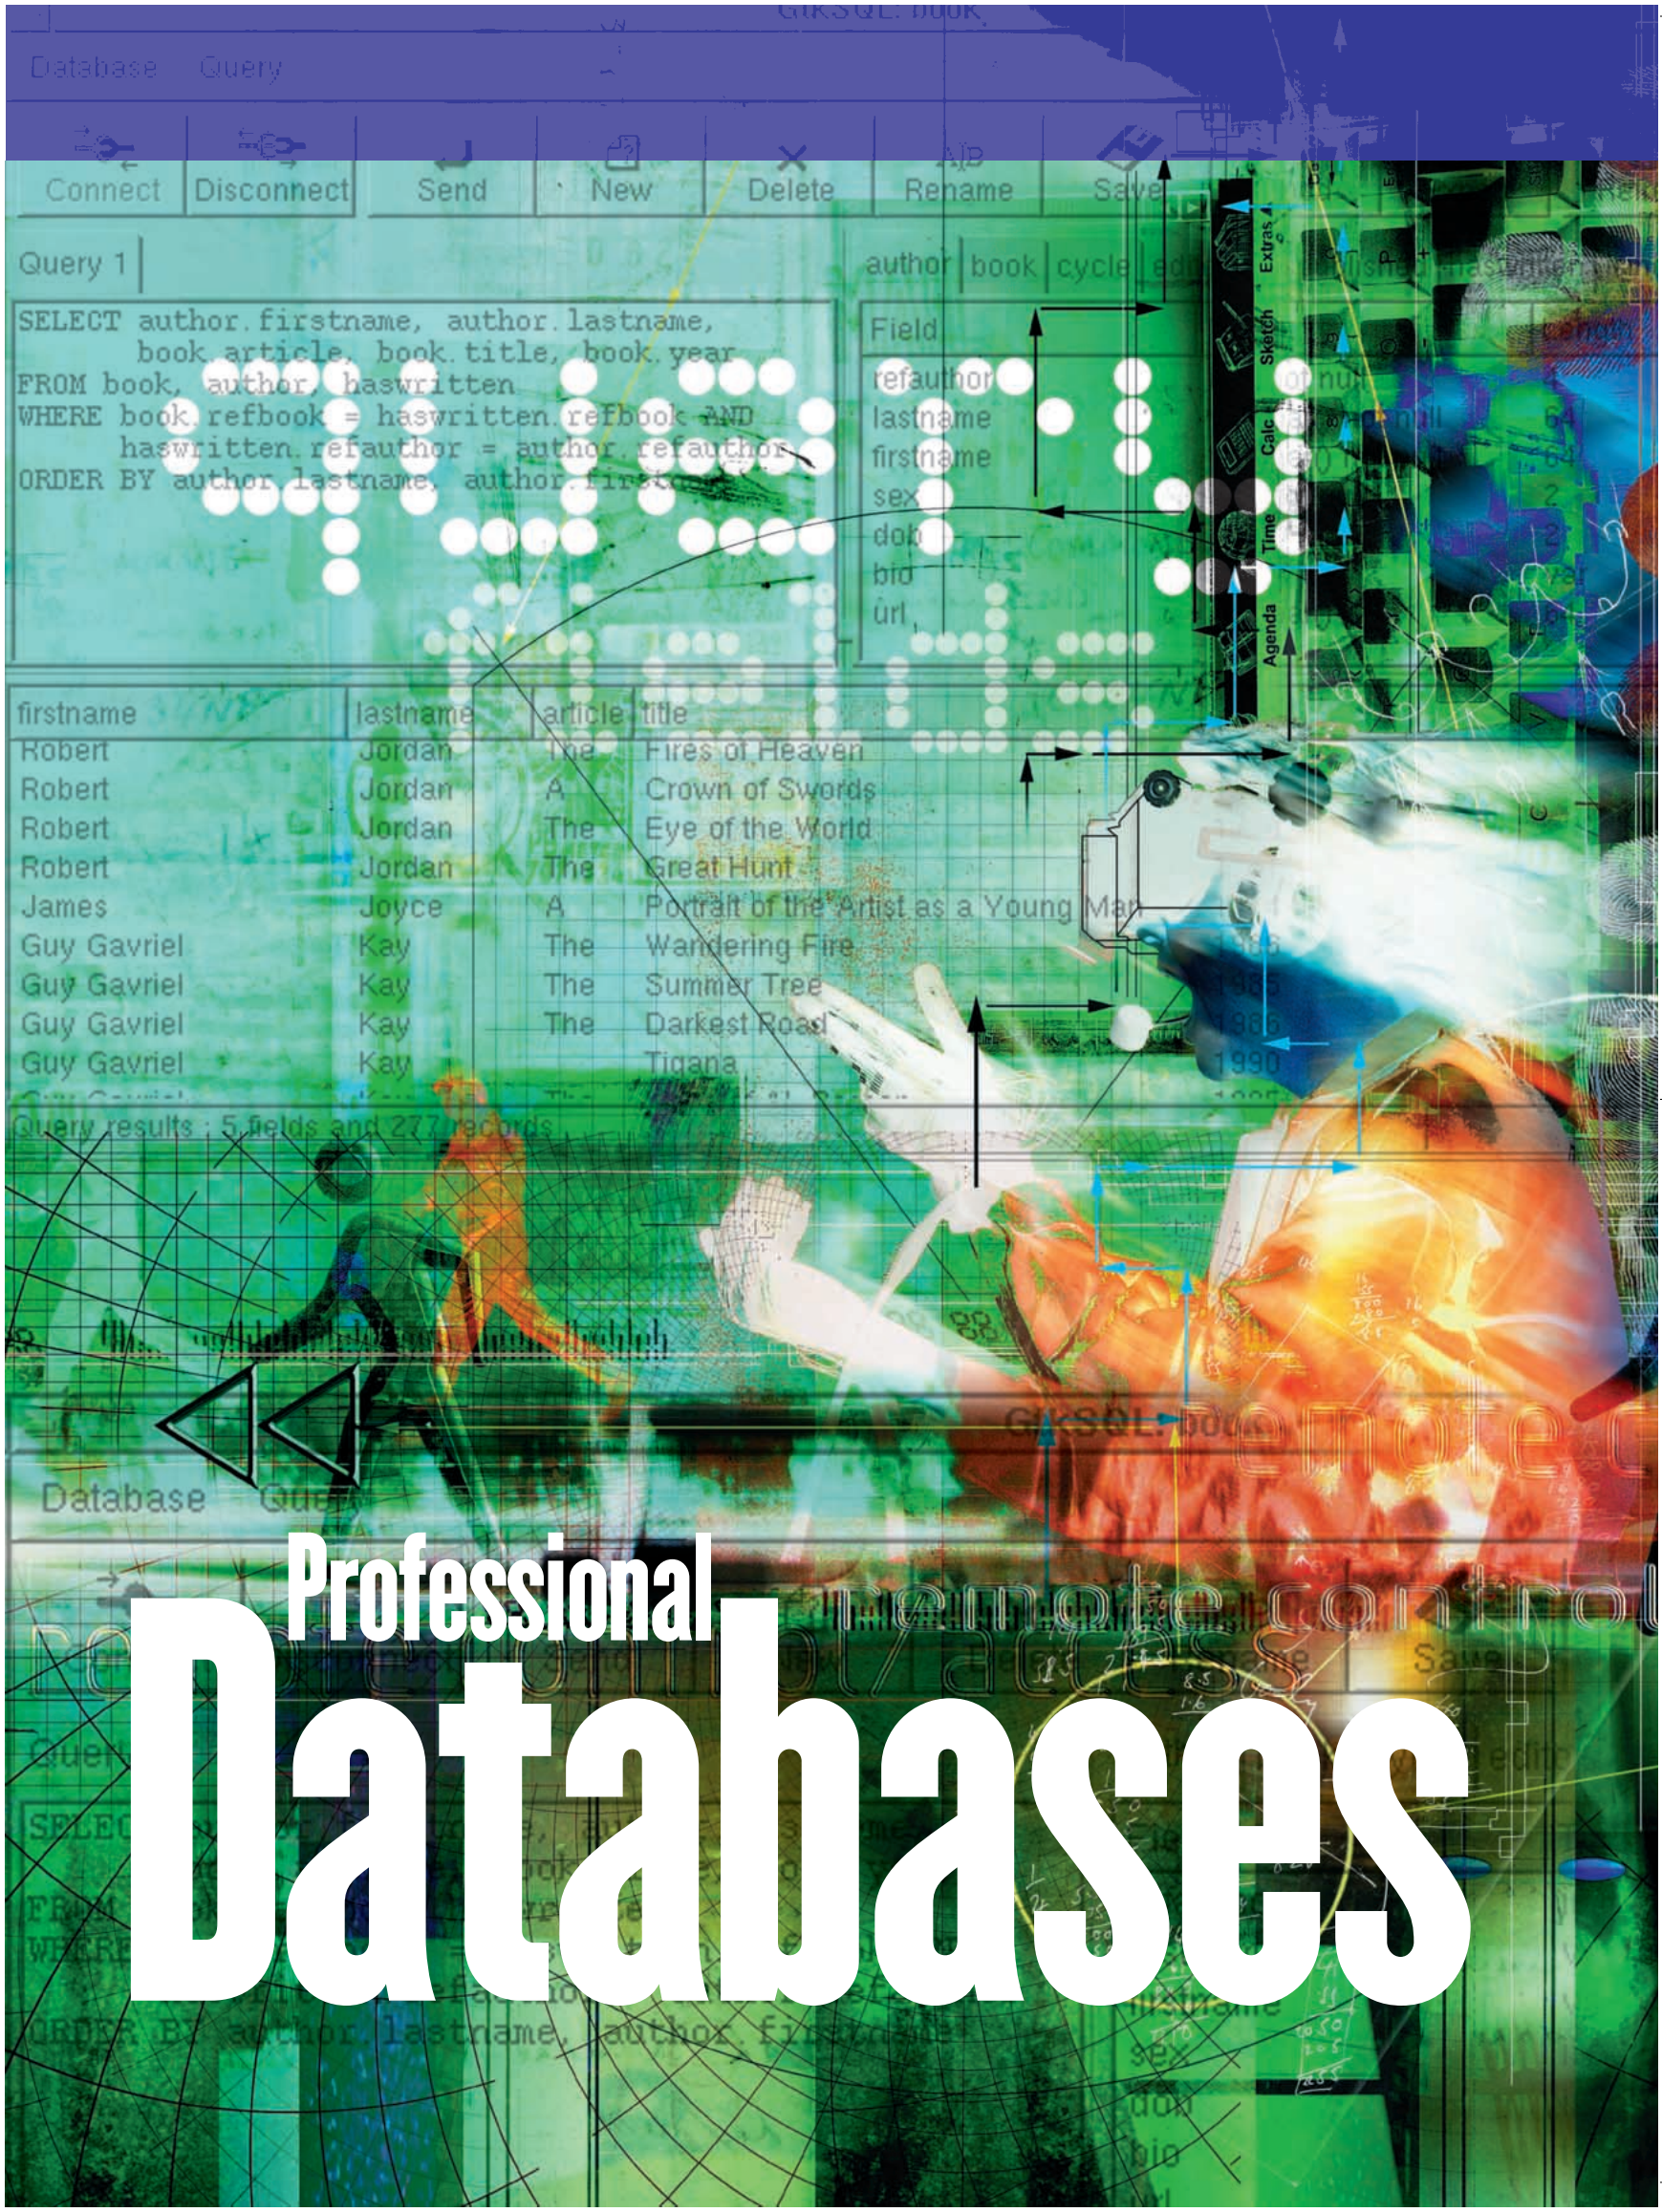

Professional Databases

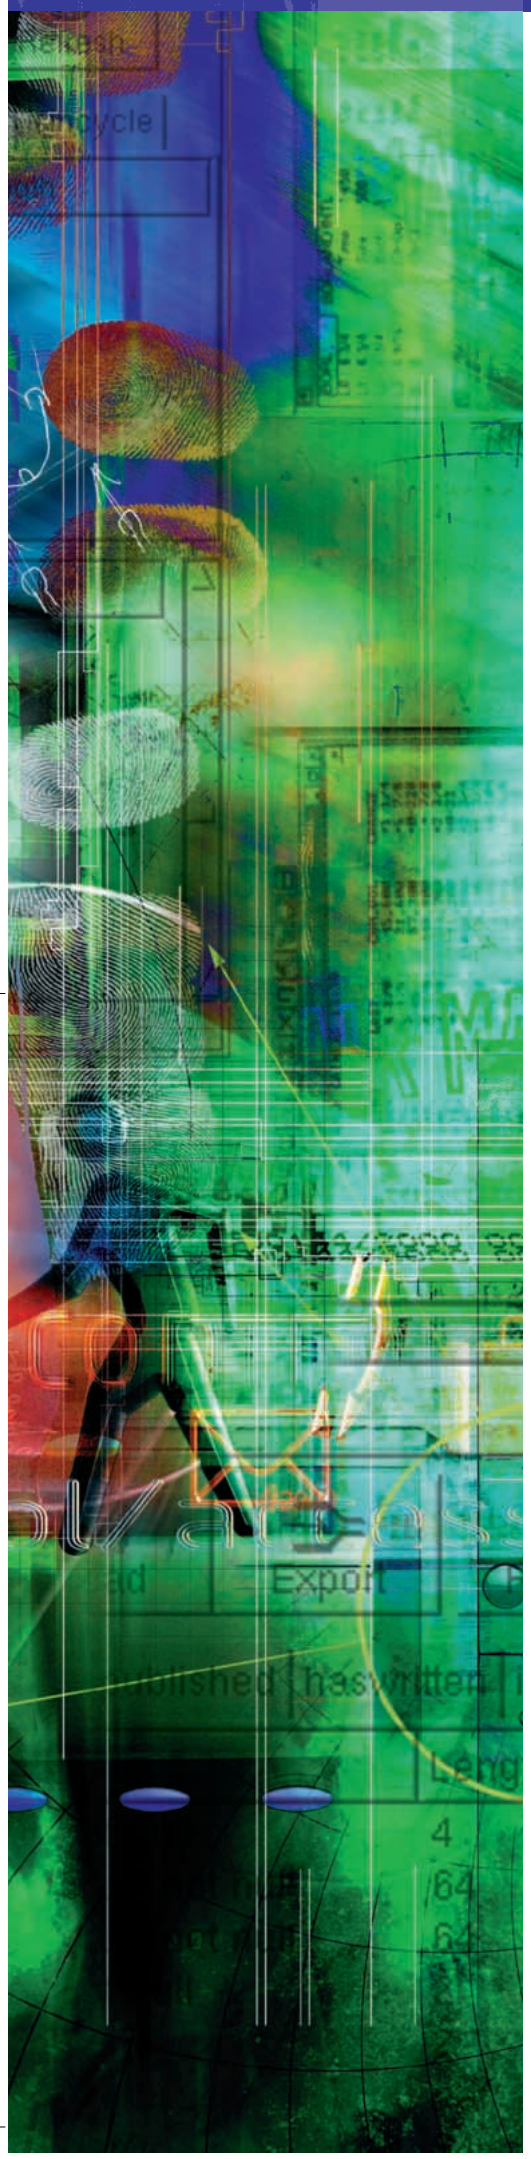

All your database are belong to us! Charlie Stross shows how best to serve up data for the masses.

A database is useful in direct proportion to the amount of data it holds. Data gets into a database because people put it there. The more people, the better: this is why databases are the classic 'killer app' for networking in business.

Sales of Ashton-Tate's *dBase* drove sales in the late 1970's and early 1980's; then heavyweight relational databases like IBM's *DB2* and Oracle built the server and mainframe market.

With Linux increasingly visible to the industry, it's no surprise that the usual suspects are all available in Linux editions; you can buy *Oracle*, *DB2*, *CA OpenIngres*, *Adabas D*, *Informix*, and others that are familiar from Windows NT and the corporate UNIX world.

However, there's a new phenomenon in the database world. A number of open source database projects exist; either as outgrowths of earlier academic research (*PostgreSQL*), developments of lightweight storage tools for websites (*MySQL*), or open-sourced commercial products (*InterBase*). The success of these products has been intimately entwined with the Linux platform, and as Linux has spread, so the open source developer community has grown. All of these databases now have commercial support and development offerings to rival the professional products, and unless you plan to run a multi-terabyte data warehousing operation they can handle just about anything you care to throw at them.

What is a database?

A database is a repository for structured information, configured so that it is possible to retrieve it easily and reliably. This definition begs so many questions that an entire field of computer science had to be established to clarify it. A filesystem is a database for storing and managing files, but it lacks most of the features we commonly associate with a database program.

Computerised databases first became practical in the 1960's, when large disk drives and tape-based file storage systems allowed significant quantities of information to be stored. One of the most important features of a database was that the information in it had to have a recurring, regular structure – for example, driving license records might contain the drivers' names, fields for address information, age, license expiry date, and codes for the vehicle categories they're licensed to drive and any penalty endorsements on their license. A second feature is that the data is organised to make it easy to search for records that match some desired characteristics (for example, all drivers called "Fred" aged 36 or over, living in Swindon, with six or more penalty points).

A brief history

At first, databases were designed around fixed data structures and the programmers who developed them wrote their own disk access libraries for storing and retrieving data. Later, libraries of software for handling low level storage and indexing became available. The problem with this approach was that changing the format of the data records was extremely difficult, as was adding new search routines: while it's fine for simple applications such as an address book, it is difficult to use such a database to represent, for example, an organisation's infrastructure.

During the 1970's, IBM researchers developed the relational model: a formal mathematical specification of how a database can be arranged that maximises the flexibility with which data can be inserted and retrieved. The relational model takes an abstract view: each database consists of a set of tables. A table consists of columns (named data types), and a row in the table is a record (consisting of a group of values in different columns that are associated with each other). A key is a column that contains unique values, so that no two rows have the same key. Identical keys can, however, be used to associate rows in different tables, as long as they contain a common key field.

A formal logical framework for performing operations on these tabular data structures was then developed. SQL, Structured Query Language (the standard tool for updating/querying relational databases) is essentially a relational algebra system implemented as a declarative programming language.

While relational databases such as *DB2*, *Oracle*, or the open source alternatives are extremely flexible, small tasks such as the web address book or password database may be better implemented using a low-level library. A variety of these toolkits are available on Linux.

Non-relational databases

You can implement a non-relational database using any programming language that lets you read and write the contents of files, However, not all ways of accessing data in files are equally efficient. Relational databases typically store their tables using specialised on-disk data structures such as hash tables; if you know the design of your database in advance, »

"A database is a repository for structured information, configured so that it is possible to retrieve easily"

Further reading

For an overview of the relational algebra, see *An Introduction to Database Systems*, 6th edition, by C. J. Date (Addison-Wesley, ISBN 0-201-82458-2).

« you too can use low-level storage libraries. The advantage of doing so is that they're faster and more efficient than a relational database engine which has to figure out where the abstract data you're requesting is stored before it can retrieve it and hand it to you.

The most common simple database libraries on UNIX (and Linux) are the dbm family of routines. Provided with early UNIX systems, the current representative of this family is *gdbm*, the GNU database manager. *Gdbm* is a library of routines that manage database files containing key/data pairs: you can store, retrieve or delete data records by key, and you can traverse the records in the database (in an unsorted sequence). Each *gdbm* file can have only one "writer" process at any time, but supports many concurrent "reader" processes. What you store in the data records (which are limited to 4K) is up to you, but typical uses include saving password databases: the key might be a user's unique ID on a system, with the record containing the user's details – including full name and an encrypted password – in a form that can be saved from or loaded into a C structure.

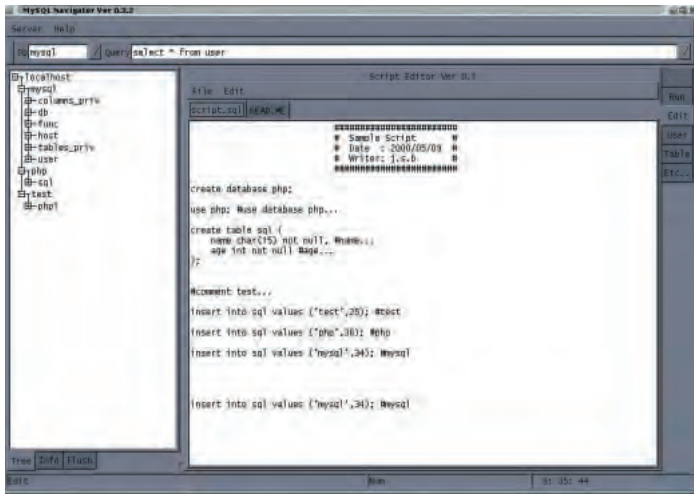

***MySQL Navigator* is a more comprehensive tool for managing MySQL databases. Here it is being used to edit a script for creating a table.**

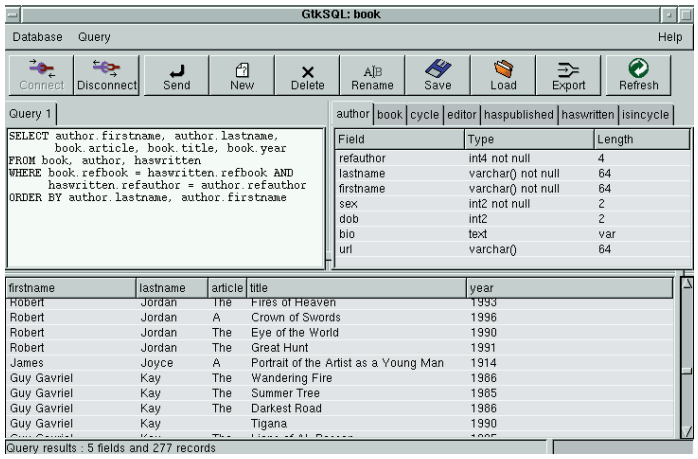

***GtkSQL* is a Gtk-based query tool for *PostgreSQL*. Its main job is to make it easy to compose and execute SQL queries on an existing database.**

Building blocks

Gdbm is a traditional non-relational library, but there's another, more complex one that is prevalent on modern Linux systems: *Berkeley DB*. *DB* is a family of groups of functions that provides record-oriented file access with transactions and a modular interface. Berkeley *DB* includes b+tree, queue, extended linear hashing, fixed, and variable-length record access methods, transactions, locking, logging, shared memory caching and database recovery. It's thread-safe and supports C, C++, Java, and Perl APIs. It also provides for multiple simultaneous readers and writers, and guarantees that all changes are recoverable, even in the case of a catastrophic hardware failure during a database update. What it doesn't provide is end-user interfaces, data entry GUI's, SQL or ODBC support, or other standard database interfaces. What you get are programmatic building blocks that allow you to easily embed database-style functionality and support into other objects or interfaces (see www.sleepycat.com for details).

If this sounds a bit low-level for you, there are other non-SQL based database engines available for Linux. Possibly the most interesting is *Harbour* – a clone of *Clipper*, a compiler for the dBase III+ language. *dBase* was the standard database system used on PC's in the 1980's and early 1990's. Although it was sold as a relational database, it was relational only insofar as it was possible to open multiple databases and carry out joins on their contents by hand. In reality, *dBase* provided a high-level procedural language for managing indexed files of records, with a number of useful additional functions for painting screens, printing reports and doing useful tasks. *Clipper*, a clone of the original Ashton-Tate *dBase*, was a compiler for dBase language programs.

While *dBase* was displaced by true SQL-based relational database engines, it continues to have a following as it's useful when you want to build a database rapidly and know the structure of the records you'll be working with.

Harbour (from www.harbour-project.org) is a GPL'd *Clipper*-compatible *dBase* compiler. Still under development, the goal of *Harbour* is to provide an open source *Clipper 5.2* compiler that produces native binaries for a variety of platforms (including Linux and Windows). If you have written *dBase* applications in the past, want to maintain an existing application or port it to a new platform, or just want a lightweight tool for managing databases without the complexity of the full relational model provided by the big SQL-based systems, *Harbour* may be exactly what you need.

Next generation

dBase was arguably the first fourth generation language (4GL) – a programming environment designed to make writing business applications software easy. Other 4GLs have been ported to Linux, but these are mostly commercial systems. For example, *Informix 4GL* – a relational database system with its own high-level language environment – is available on Linux as a commercial product. More recently, Borland's Delphi language has been ported to Linux in the form of *Kylix*. *Kylix* provides an integrated development environment and compiler for the Delphi language (basically Borland's own descendant of Object Pascal, with added data handling stuff), along with a comprehensive class library and a

user interface design tool. *Kylix* may well be a big hit with Windows programmers seeking to port applications to Linux – the ability to simply recompile and have a working graphical application is a big win.

This is only scratching the surface of database systems on Linux. You can find an exhaustive list at: **http://linas.org/linux/db.html.**

Finally, there is an open source COBOL implementation for UNIX and Linux called TinyCOBOL. You'll find it at **http://tiny-cobol.sourceforge.net/**. It implements a subset of COBOL 85, and is best passed over with a shudder.

Relational databases

All serious databases are client/server applications: a faceless server program rides herd on your data and handles queries, while multiple users send queries or data to it using client programs. The server's job is to handle these requests and maintain consistency in the data, which is stored either in files or a raw partition on the hard disk.

Consequently, what you tend to find is a core server daemon (such as mysqld) which is usually started when the host machine boots, and runs continuously unless someone orders it to shut down. Administrative commands – such as shutting down the server, telling it to flush its buffers, collecting data on how it's performing, create a backup copy of a database, and so on – are sent to the server from a special administration tool (such as *mysqladmin*). Other commands that interact directly with the database may be issued by a query monitor.

"Most databases offer a variety of sophisticated graphical tools for designing data structures."

A query monitor (such as the mysql program) is a command-line shell that lets you interactively type SQL queries. These are sent to the server daemon which executes them and returns the results to the monitor. While these are useful – and, by letting you issue SQL commands from shell scripts, provide a lowest common denominator for database integration with other software – their user interface is extremely primitive. Most databases offer a variety of sophisticated graphical tools for composing queries, formatting reports, and designing data structures for creation using the database's data definition sublanguage.

Databases are often used for information storage by other applications. To this end, they come with well-defined APIs in a variety of languages: the APIs generally let a programmer write SQL commands and send them to a database server, and get back usable results, so that a C program can call on the services of a database when necessary.

Different approach

One point to note is that the SQL language is a moveable feast. While almost all databases implement a core of it, most add their own extensions and in some cases the standards body has been nobbled to specify that one vendor's way of doing something is the right way to do that. An example is *PostgreSQL* and outer joins – it's technically not compliant with SQL92 because it can't do them using the ANSI syntax, but it's functionally equivalent, because it can get the same results using a different approach. The upshot of this is that SQL code may need to be re-written when moving from one database server to another.

The server daemon itself has a number of sub-components, but these are usually hidden from the user. SQL is a complex language, and all relational databases have

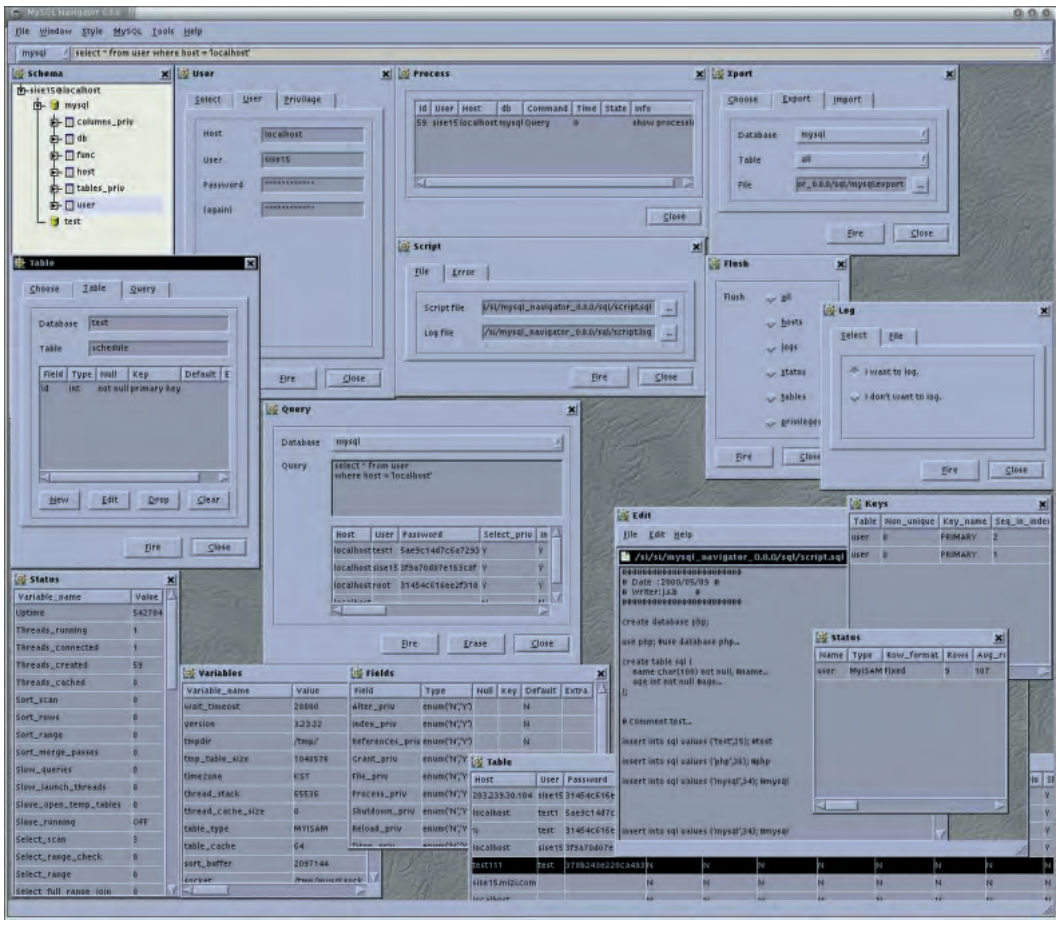

***MySQL navigator* running under the Qt graphical environment. Note the ability to examine and edit *MySQL's* control variables, change the internal storage format of a database, control logging options, and generally manage a *MySQL* server. This isn't a tool for users, but it makes your life as database administrator a lot easier.**

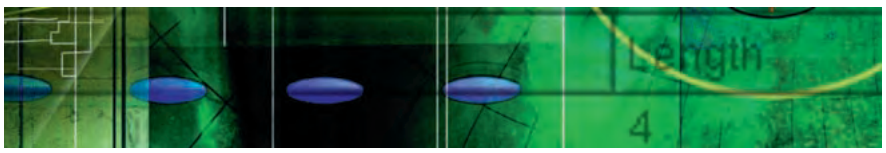

« optimising compilers built in that attempt to re-order a SQL query for maximum efficiency of execution. The server may be multi-threaded or fork child processes to carry out each client request, but in either case it needs a mechanism for maintaining control of its internal state. It is vital to prevent one UPDATE statement from overwriting a table that's in use by another client's SELECT statement, for example. This gives rise to the concept of a transaction: a group of operations that are carried out in such a way that the database's internal state is always consistent – either they have all been executed and finalised, or the database can return to the state it was in before the transaction commenced.

"Individual database servers are vulnerable to external problems – hard disk crashes, power failures and so on."

Individual database servers are vulnerable to external problems – hard disk crashes, power failures, burst water pipes, and so on. There has therefore been a lot of work undertaken in the field of replicated and distributed databases. A replicated database is one where copies exist in multiple locations, albeit only updated by one server at a time, and a true distributed database is one where parts of the data are stored by each machine in a cluster, and a failure of one machine cannot result in actual loss of data.

MySQL

MySQL emphasises speed over security and may not be ideal for mission critical jobs.

MySQL is the most commonly used open source relational database. Developed by Monty Widenius of TcX, *MySQL* is primarily designed for speed and light hardware demands, although it is relatively feature-rich.

MySQL was originally licensed commercially, but is now an open source product, released under the GNU General Public License (GPL). You can buy a commercial license if you want to use *MySQL* in a manner incompatible with the terms of the GPL, or if you want support from MySQL AB, the company spun off from TCX to cope with demand for the product. Support is also available from NuSphere.

The core component of the product is a multi-threaded database server (called mysqld) that speaks an extensive subset of ANSI SQL92. There are one or two variations from the language standard; in particular, sub-SELECTs don't work (e.g. SELECT * FROM table_foo WHERE id IN (SELECT id FROM table_bar)), and it doesn't support SELECT INTO TABLE.

If you find you really need a mouse-driven application, graphical shells for *MySQL* are available – notably *MySQL Navigator* (http://sql.kldp.org/mysql/). In addition, *MySQL* comes with a variety of command line tools: the mysql query monitor (which lets you run SQL statements interactively), the mysqladmin tool for administering the database server (for example for starting up and shutting it down in a controlled manner), tools for backing up and restoring databases, and maintenance utilities.

The ACID test

MySQL is not an ACID database: ACID (atomicity, consistency, isolation, and durability) are essential traits in a mission-critical system where the effect of losing data can be severe. Rather, *MySQL* places an emphasis on speed: instead of supporting transactions, it focuses on atomicity of operation. Updates are carried out in such a way that while each specific update is running, no other user can interfere with it and there will never be an automatic rollback. Atomic operations are in general far faster than transactions, and the MySQL approach to avoiding race conditions is to include simple checks before updates and to run scripts that check databases for inconsistencies and repair or warn in the event of problems. *MySQL* allows table level locking and it's possible to lock a table for consistency checking while using INSERT DELAYED to queue updates until it's safe to apply them. If you really and truly need transactions, *MySQL* provides a table type that supports COMMIT and ROLLBACK operations, subject to a performance penalty. However, if you value integrity over raw speed, *MySQL* is not ideal.

MySQL supports stored procedures – chunks of SQL that are compiled and stored in the database

MySQL support

***MySQL* is supported commercially by at least two companies: MySQL AB in Sweden (see www.mysql.com), and NuSphere LLC, in the United States. MySQL AB provide support in a range of packages including basic email support, extensive email support (including debugging simple customer problems), login support (an engineer will log into your server and try to diagnose/fix problems for you), and extended login support (an engineer will try to optimise your system to your satisfaction).**

NuSphere take a different approach. They sell packaged MySQL distributions, including printed manual, O'Reilly pocket guides (for some of the other software supplied with the package), a CD-ROM with copies of the software optimised for a selection of different configurations, and a useful web-based administration tool. NuSphere aren't simply making money off of MySQL AB's work, however; they're actively contributing to the MySQL code base, for example by adding facilities such as row-level locking.

server, so that clients can invoke a stored procedure by name instead of re-compiling the whole thing every time they need it. However, triggers (stored procedures which are triggered whenever a particular event occurs) aren't available yet, but are due in a future release.

You can communicate with the *MySQL* server via TCP/IP or UNIX domain sockets. In fact, support for web applications is one of the application's main jobs. Consequently there's an extensive access control mechanism built in. User privileges are stored in the mysql master database (which contains all the server configuration information), and users can be granted authority to select, insert, update, delete, alter, index, create or drop tables, on a per-table basis.

Other privileges (including the ability to grant privileges to other people) are available. Users are defined in the mysql.user table, which stores an encrypted password for each user and lists the hosts from which they're allowed to connect to the server. Each connection to a *MySQL* server is subject to vetting based on your identity and the host you're connecting from. Warning: a vanilla *MySQL* system is not secure! If you install one, you should read and understand chapter 6 of the manual (which describes how to secure a *MySQL* installation) first.

Distributed data

MySQL doesn't support fully distributed operations, but it is moving in that direction. Since about release 3.20, the database has supported logging; with this switched on all transactions, or transactions which modify tables, are logged to a text file somewhere in the filesystem that *MySQL* runs on. This means that you can take a backup of a database and a logfile, and by applying the logfile to the backup you can bring it up to date – especially handy if you are using something like *rsync* to copy the logfile to a remote system.

Replication support is provided by way of the logfile. You can tell a *MySQL* server that it is one partner in a master/slave relationship with another server. The slave will periodically poll the master, suck up the binary logfile, and play it back against its own database until it's in step. You can use this to build fall-over capabilities (the master stumbles, but your client software is smart enough to try to establish a fallback connection to the slave server), or for clustering where a single master updates the database, but a pool of slaves is available to carry out queries. While it's possible to define a set of master/slave relationships that run in a circle, there's some latency in the time it takes for updates to propagate between servers – it is not safe to assume that the answer you're getting from a query to one server is the same as the answer that all the others will give. Promised for a future release is the ability to log the machine an update came in from, so that true clustering is possible – that is, updates will go to all the hosts that need them.

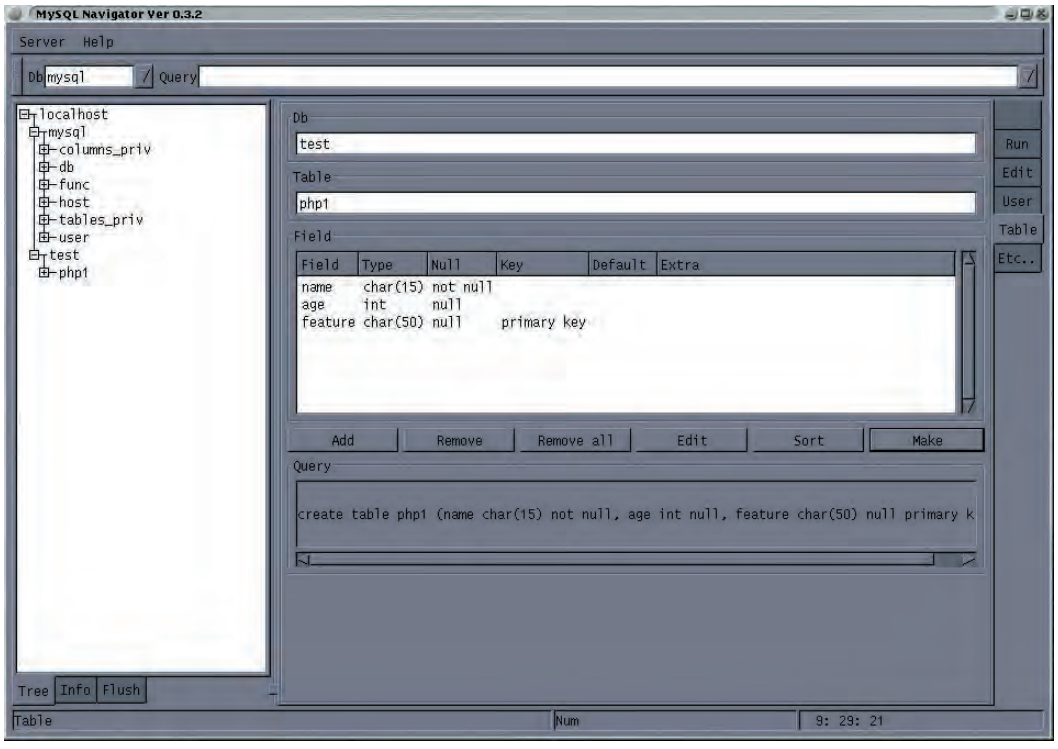

MySQL Navigator **can provide a hierarchical view of databases and tables on a *MySQL* server.**

The *MySQL* manual is exhaustive (it runs to 650 pages when printed), and in addition to a complete language reference it has extensive notes on performance tuning, a lengthy tutorial, notes on performance-critical issues such as table types (it provides ISAM, MERGE ISAM, HEAP, and Berkeley DB tables) and how to administer a server. Among the tools that come in the box are a bewildering variety of APIs (for C, C++, Eiffel, Java, Perl, PHP, and Python), benchmark tests, tools, and extensive documentation.

"Don't expect to spend less than a few days learning your way around."

Setting up a system based on *MySQL* is a large job, so don't expect to spend less than a few days learning your way around any relational database system (Remember, there is a reason why 'Oracle DBA' is a job description!).

PostgreSQL

PostgreSQL's ACID features make is a safe bet where data integrity is paramount.

PostgreSQL, now at version 7.0.2, is an open source database server. Its primary emphasis is on functionality rather than speed, although the current version is at least as fast as *MySQL*. *PostgreSQL* is descended from the original Ingres research database, which was commercialised by Ingres Corporation. The Ingres research project spawned an object-relational database server called *Postgres*. This in turn spun off Informix Universal Server, then the *PostgreSQL* project in 1996, with a bunch of open source developers who were acutely aware of the need for a heavyweight open source relational database system. They started by fixing bugs, then worked on adding new features and improving performance, and the result is *PostgreSQL* 7.

PostgreSQL support

PostgreSQL is supported commercially by a company called Great Bridge (www.greatbridge.com). They'll sell you a packaged, QA-tested distribution of PostgreSQL, complete with printed documentation. They also sell technical support packages and professional consulting services. These services are not cheap but their high-end package provides 24x7 coverage, unlimited phone and email incident support, and a team of (named) support analysts assigned to the customer. Unless you desperately need a distributed database, PostgreSQL provides similar functionality to Oracle – and you can buy a lot of support for the price of a multi-processor Oracle license.

PostgreSQL uses a client/server architecture. A client application that wants to use the database connects via TCP/IP to the Postmaster master server, and this spawns a *PostgreSQL* server to handle the client's session.

PostgreSQL supports almost all of SQL92 and also has object-relational features. A table (also referred to as a class) has numerous rows (instances), and you can operate on it using SQL commands, either interactively via the pgsql query monitor, or via a language API. Tables can inherit attributes – you can define a child table that inherits all the column names of a parent class (table) and then adds more of its own. SELECTs can be applied either to a specific table, or to a table and all its descendants. In addition to providing the normal range of types (INT, VARCHAR, FLOAT and so on), *PostgreSQL* provides a rich selection of types, in particular non-atomic values. Columns can contain arrays of items, for example multi-dimensional arrays of integers, and you can access arbitrary slices of the arrays by specifying both lower and upper subscript bounds. There are other categories suitable for geographical databases, such as geometric types (line segment, point and polygon coordinates); network data types (such as ethernet MAC address and IP network address); and temporal data types (date/time/interval). You can convert between datatypes using CAST, and a range of functions are provided for massaging each type.

PostgreSQL's data manipulation language is far richer than *MySQL's*. You can give table names aliases, permitting self-join operations (where a join is performed on two instances of a table). It also support Sub-selects and SELECT statements can be chained together using UNION, EXCEPT and INTERSECT clauses. The one chunk of SQL92 that it doesn't fully support is outer joins, but this can be simulated using subqueries.

More ACID

The big difference between *PostgreSQL* and *MySQL* is that *PostgreSQL* aims to provide ACIDity – Atomicity, Consistency, Isolation, and Durability. The database is required to never lose data – no matter if someone trips over the server's power cable, or a disk fills up in the middle of a transaction. Anything that gets committed to the database must be retrievable, even if this entails a performance hit. *PostgreSQL* provides transaction support, including serialisable isolation (a multi-statement transaction in process won't see any updates to the database happening due to other processes). Locking is provided at table and row level, and unlike *Oracle*, *PostgreSQL* doesn't have raw access to a disk partition – so when processing transactions it calls **sync()** after each update. This may slow things down considerably, depending on your system, but it can be switched off (if you're willing to assume your filesystem is robust enough to save your bacon in event of trouble). It is a good idea to run a mission-critical *PostgreSQL* database on top of a journalling filesystem.

Security is assisted by a complex domain-based authorisation system. It's possible to configure *PostgreSQL* to tunnel its TCP/IP connections using SSH – this is essential if you're moving confidential material over a public network. *PostgreSQL* also supports Kerberos authentication.

PostgreSQL provides stored procedures and triggers. Stored procedures can be written in several languages, including C, PL/TCL, PL/Perl, and PL/PGSQL, rather than just SQL. A trigger calls a stored procedure each time a row is modified in the target table. A different type of trigger called a rule allows SQL queries to be executed whenever a table is accessed.

The standard way of interacting with *PostgreSQL* is via a SQL query monitor (pgsql) or from an application written using a database access library. However, there's a nice, useable GUI supplied with the distribution called *pgaccess*.

PostgreSQL has been saddled for years with a reputation for slowness, which is a hangover from its past as an academic research project – and backed up by tests on versions up to 6.5.3. The latest version is different: it performs well under heavy load, and runs AS3AP and TPC-C benchmarks slightly faster than *Oracle* and only slightly slower than *MySQL*. LXF

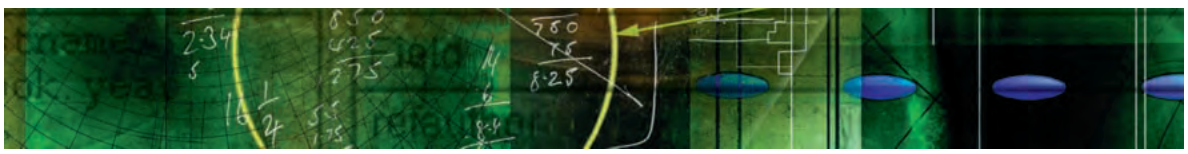

LINUXformat Desktop

All you need for configuring, using and customising your desktop

THIS MONTH...

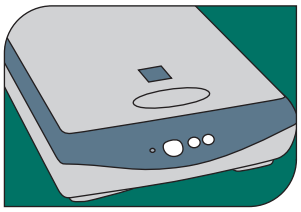

NEWS

Printing made easier

Turboprint for Linux

German firm Zedonet have unveiled a new printing system for Linux. Until the advent of printing systems such as CUPS (see issue 10) Linux was pretty much in the dark ages when it came to using modern desktop printers – particularly inkjets with their range of printing modes and colour combinations.

TurboPrint claims to build on the existing Linux printer technology (lpr and ghostscript), but provide more intelligent and specific printer drivers for a range of popular inksquirters. A huge range of printers from Canon, Epson and Hewlett Packard are available, with more on the way.

Zedonet are releasing a free download version of *Turboprint* (with some restrictions), with the commercial version being priced at 19.95 Euros.

The *TurboPrint* software was previously available for the Amiga platform, where its custom drivers and excellent printing quality earned it a huge following.

http://www.turboprint.de

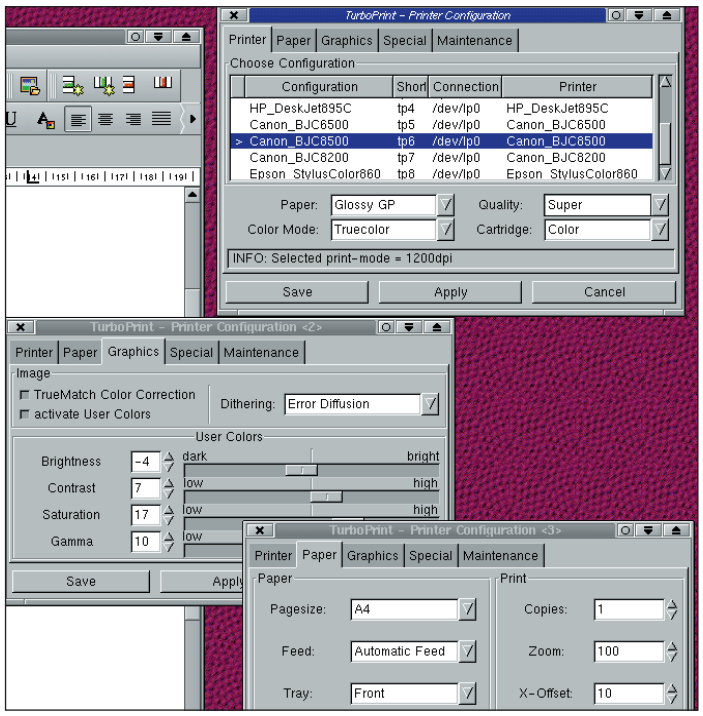

Turboprint supports a wide range of popular inkjet printers.

Linmodems update

Motorola support Linux modems

At long last, you may think, Motorola has announced limited support for their SM56 series of data/fax modem cards. The new drivers are claimed to be fully functional, and will drive both the SM56 PCII and SM56 AC-L devices. In an further encouraging move, Motorola are offering email support for Linux users of these devices, and have an ongoing support and development program for the drivers, which can be downloaded from www.motorola.com.

PCI modems are usually a problem for Linux users, as they rely on interoperation with the CPU to function properly – the host computer takes care of various processor intensive signal processing tasks which is costly to build in to the modem itself, a technology known as Host Signal Processing (HSP). On the Windows platform, this relies on the manufacturers providing a driver which takes care of this necessary function, but as you may guess, few manufacturers bother to come up with similar drivers for other OSes. Becausee of the custom nature of such drivers, and the reluctance of the manufacturers to give out details on exactly how they work, it's an almost impossible task for third-party open source developers to create drivers.

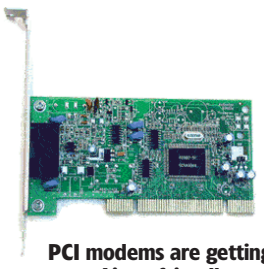

PCI modems are getting more Linux-friendly.

Lucent are the only other major modem manufacturer who have produced support for some of the Winmodems.

Office suite snapshot

OpenOffice update

A new snapshot of *OpenOffice*, 627 has just been released. Although by no means even nearly complete, this project has come a long way since the initial decision by SUN to open up the source code.

While *StarOffice* was a functional office suite, it depended heavily on other commercial code for a number of functions, so the open-source effort has concentrated on rebuilding the framework of the whole application, and writing new code to replace the functional, but proprietary code which has been lost in the transition.

The *OpenOffice* team still consider this latest snapshot to be Alpha software, but as with many projects, they are depending on bug reports from the community as a whole. If you'd like to help with this effort, you can download source code or binaries from the project website, and we'll try to fit it all onto our coverdisc next issue too.

http://www.openoffice.org

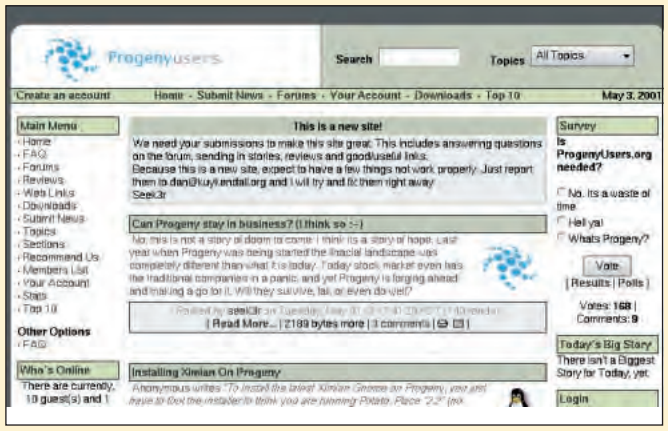

www.progenyusers.org has a familiar PHP look to it.

Distro support

Progeny Users online

Progeny users now have their very own website – and it's proving more than a little popular. Although only up and running for a couple of days at the time of writing, the site has already received over 6500 views from Progeny fans and curious visitors.

Running on PHP-Nuke (the same web management system employed on our own *Linux Format* site), the site offers help and information about Progeny, useful links and a selection of web forums. If you're a Progeny user, you should check it out at:www.progenyusers.org

Progeny Linux is a commercial version of the popular and reliable Debian Linux distribution which offers printed documentation, support and many other features you could expect from a commercial distro. We have a full review on page 28 of this issue.

http://www.progenyusers.org

Boot up and go?

The desktop is the big market, but is Linux ready?

In order to capture a wider slice of the Desktop market, Linux needs to adopt an appliance-like mentality. The term "computer user" doesn't do justice to the wide range of people who use computers, and the purposes they use them for. And people who use them as a tool to do a specific task far outnumber those who engage in computing as a hobby, or a pastime, or who are even remotely interested in what their computer does.

Let me diverge into an analogy, and talk about the two-wheeled vehicle industry.

People who want a cheap and convenient way to commute to work (treating their vehicle as a transport appliance) choose a Spacey or a Cub90 or whatever the latest pseudo scooter is. People who actually want to enjoy the experience of using their vehicle, who take an interest in it for itself, are more likely to choose a Ninja, Bandit, Guzzi, Trident or Sportster, depending on their individual feelings on style.

Like the motorcycle industry, there is no reason why the world of computers can't offer appliance solutions for computing. In fact it already does – Apple have been trying to turn computers into appliances for years, and have achieved quite a success with the iMac.

The great thing about Linux is, that it's possible to have both in the same operating system, though perhaps not the same distribution. Efforts by distributions such as MaXos, Corel, and even Caldera's eDesktop have attempted to deliver trouble-free, easy to install and use solutions for Linux, with varying degrees of success. Surprisingly perhaps, Corel's efforts were the most ambitious, but also one of the least successful.

There doesn't seem too much money (yet) in the idea of producing distros for home desktop use, but it may not be too long before people wake up to the fact that for businesses, a viable Linux desktop solution could be even more attractive than just using Linux for their web or mail servers.

Nick Veitch
Is the Editor of *Linux Format*, and runs Linux on every desktop machine he owns.

TURN-BASED STRATEGY

Jagged Alliance 2

■ **DEVELOPERS** Tribsoft ■ **PRICE** 59DM ■ **WEB** http://www.tribsoft.com

Hiring mercenaries, training militia, freeing a small nation from the yoke of dictatorship – it's all in a day's work for Paul Cavanagh. But, Oh how he wishes it wasn't.

When this game turned up in the office in its very smart DVD style case, we were all a bit shocked. Nobody told us that *Jagged Alliance 2* was under development, it hadn't come from prime Linux developers Loki, and best of all it isn't a first person shooter. While fps games are great, it is possible to have too much of a good thing, and variety being spicy, well, we were quite chuffed.

Post-installation, things started to look rather less rosy. Cast your eyes over the pictures gracing this page, and you'll begin to understand. Here's a game that can be compared (unfavourably) with such Amiga classics as *Syndicate*, *Cannon Fodder* and the more recent *Wasted Dreams*. Those games were pretty much all top-notch on low-spec machines – but here we are with a minimum spec of a Pentium 2 at 233Mhz with a graphics card and nothing has improved. The whole thing looks a little home-made with poor animation, slow scrolling and clumsy characterisation. But, come now, looks aren't everything – underneath this archaic surface may lurk an intriguing, addictive little beauty. Just maybe.

The game's plot is cleverly revealed through a laptop computer that can be accessed through the map screen. Here you can view the history of the small country that you have been asked to liberate, and then go on to hire mercenaries. Through the laptop you can access email, surf the web (where you can buy weapons and visit mercenary agencies' homepages) and review your finances. This part of the interface works very well and adds a little credibility to the proceedings – after all it seems perfectly believable that small armies could organise themselves in this way.

Once your mercs have been dropped off you are immediately engaged in battle. This is where things really start to fall apart. At first I found that it didn't seem to matter what tactics I used, I would always be at a major disadvantage, with my troops failing to hit the enemy, while taking numerous hits on every turn. After lengthy consultation with the manual I discovered how to get my troops to aim. Ah-ha! Even with this skill, my poor troops still suffered terribly. It is quite possible to be standing right next to an enemy soldier when you unleash a torrent of gunfire and not hit them once. Not very realistic. Or fair. Apart from this flaw, the turn-based combat system is laborious and not terribly exciting. It occurs to me that turn-based combat has its origins in pen, paper and dice role-playing games, so why do we still have it in computer games? Although there are some great turn-based games (UFO anyone?), real-time action is at least a bit more exciting in my opinion. Besides, in a battle situation, generals don't say "Right! It's your turn now! I'll just sit here and twiddle my handlebar moustache while you make your move." It just doesn't happen.

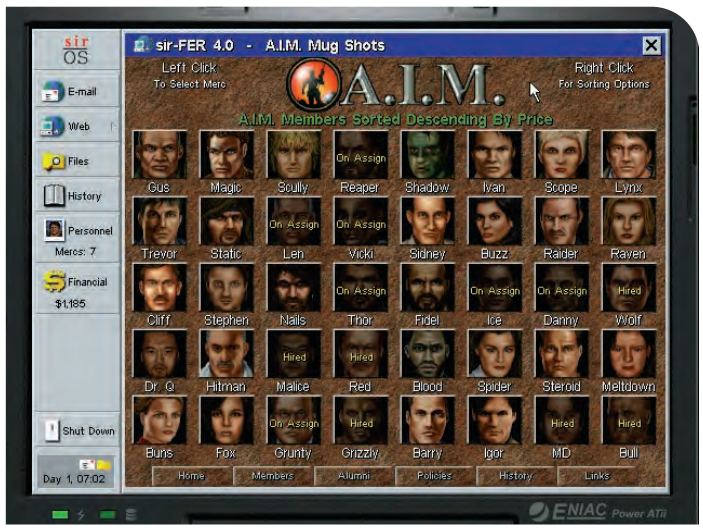

Here's the laptop display, currently using its web browser to view mercenary profiles.

Snazzy graphics, huh? This, believe it or not, is 21st Century gaming.

Thankfully there are some good bits to the game. Despite getting through a worrying number of squaddies, I did manage to make some progress. After pacifying a town I was able to train the locals to form a militia. Whenever a town with a militia is invaded the locals will automatically defend themselves. This means that you can have literally hundreds of people on your side, a nice touch.

Here's the worst bit though. Bugs. And not the Starship Troopers kind. Occasionally during a battle the game will display a little stopwatch for anything up to fifteen minutes. This is not a little irritating, but it isn't the worst. Using the time compression system (essential) can cause the game to shut down and throw you back to your desktop. This will keep happening when the in-game clock reaches a certain time, and the only way around it is to go back to a previous saved game and replay a large chunk of the campaign. And finally – drum-roll please – just as I was getting into capturing a hospital, I left the game for a while. When my screen-saver activated, it completely froze the entire game. I waited hours for disk activity to cease, and in the end, after trying to kill all the processes, I was forced to switch my computer off. I now have a completely unusable Linux partition and will probably have to reinstall everything. Be warned: if you must buy this game – and I must advise you not to – switch off your screensaver before playing. And expect to be exasperated. LXF

Linux Format VERDICT

Ease Of Use	5/10
Features	3/10
Gameplay	3/10
Value for money	2/10

An old-fashioned, deeply flawed product that really should be avoided.

LinuxFormat**Rating** 3/10

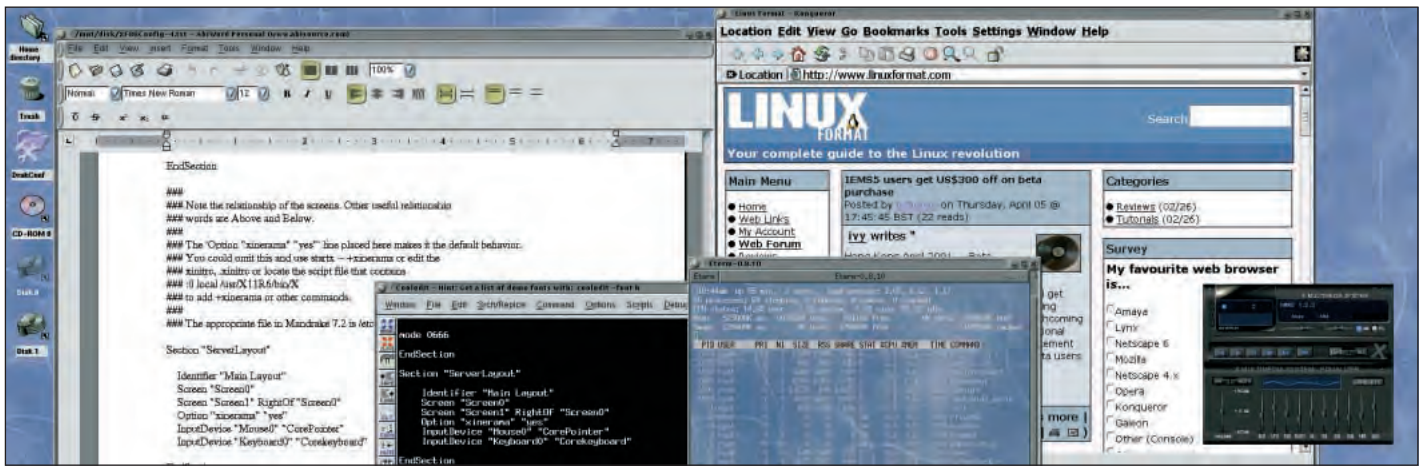

GET MORE X

Multi-Head Displays

Hoyt Duff discovers a use for that old video card and monitor left in the closet after your last upgrade.

Ever since last summer's release of XFree86 version 4, there has been much talk about one of its most intriguing features, but little written about it. In previous issues of *Linux Format*, we have titillated you with screenshots of a panoramic desktop. Now, we set the record straight and show you how to configure your hardware and edit your config files allowing you to use multiple monitors at one time, a situation called multi-head or, in its most common incarnation, dual-head.

Multi-head is nothing new: Macintosh had it in the late 1980's and Windows 98 provided it to MS users. Of course, prior to version 4, XFree86 was capable of running multiple servers and X has always been network aware, allowing the remote display of an X-client on any display on a network. Tools like *xmx*, *vnc*, *x2vnc* and *xnest* have made virtual multiple displays possible.

Actually, multi-head for Unix has been around awhile, but only available in commercial (and expensive) X-servers from Metro-X and XiGraphics. Until now, multi-head was out of the reach of the loyal, but poor, Linux user. Regrettably, this is one of those areas where our Mac and Windows brethren can justifiably gloat as multiple monitor support is much more mature in those operating systems, hence much easier to configure with a wider range of supported hardware. For the Linux faithful (as well as any other OS that can run XFree86 such as BSD), support for multiple monitors is now available; we just have to work a little harder at it – you should be used to that by now.

Why Bother?

You are likely to be asking yourself what justification there is to use multiple displays. Here are a few reasons:

1 One-Upsmanship

Be the first one in your LUG to demonstrate a dual-head configuration and you will achieve instant high status – it's simply that cool an achievement. Just ignore the smarmy Mac users in the group for now.

2 Maximum Screen Real Estate

All fun aside, this is a more practical reason. Look at it in terms of square inches and cost. You want a bigger desktop: there are two options: (1) purchase a 21" monitor or (2) purchase two 17" monitors and a second video card. The maths is left as an exercise for the student (meaning we didn't want to do it), but the display area of a 17" monitor is slightly more than half that of a 21" monitor while the cost is significantly less than half (We know your thoughts here: "Hey – can I buy THREE monitors?") and a decent new video card is about the price of a 17" monitor. It's possible (although not as desirable as we'll see later) to re-use that older video card and monitor you have in the closet from your last upgrade. Essentially, for the same cash outlay (or less) you can have a bigger display and some bragging rights.

The low-budget approach can be surprisingly affordable as you re-use that old PCI video card and 14" monitor.

3 Able to View More Windows at Once

This is the reason you use if you absolutely must justify the expense to your boss (or spouse or parents): you can be more productive. Multi-head allows you to have more large windows open at the same time without having to page through them, saving valuable time:

- Web developers can easily compare HTML pages side-by-side in full size browsers; create pages in an editor in one head, view them in the second head at the same time.
- Writers can have multiple documents or revisions open and simultaneously viewable.
- Programmers can view their development environment on one head and view the results of their efforts on the other.
- Sysadmins can display additional windows of monitoring programs, all visible at once.
- CAD/CAM, graphics designers, video editors and others can also benefit from multiple screens available for viewing at the same time.

We already hear the protests from hard-core Linux users: "That's what virtual desktops are for!" And they are correct, but with multi-head, all windows are viewable at the same time, on

Useful info

There are sample XF86Config files on you system located in /usr/X11R6/lib/X11/.

If you need to know which driver is being used by your video card, look in /usr/X11R6/lib/X11/Cards+

the same desktop. Additionally, you can still use virtual desktops to have multiple virtual displays.

Don't knock it...

Like many things in life, multi-head displays are something you should experience before you dismiss the idea. As with anything Linux, there are choices to be made as multi-head works in two basic ways. With the release of XFree86 version 4.x, Linux users can now enable multiple monitors in two ways: multiple displays (a unique desktop for each monitor) or a single display spread over multiple monitors (known as *xinerama*). In either case, you are required to have a separate video card and monitor for each display, or 'head', that you desire.

Multiple independent displays

When XF86Config is properly configured for multiple video cards, this is the default behaviour. You can move the mouse across all monitors. If one were to simply run startx, the selected desktop would appear in the primary display and X would likely appear bare in the other monitors. Why? Because many window managers are not multi-head aware, so they don't know to start a wm for those displays. Multi-head aware (to some extent) window managers include *Enlightenment*, *WindowMaker* and *Sawfish*. The multi-head capabilities of wms are not well documented.

To use .xinitrc to start non-multi-head aware window managers in separate heads, edit it as follows:

```
.xinitrc
twm :0.0 &
fvwm :0.1 &
exec fvm2 :0.2
```

substituting your preferred wm as appropriate.

You may need to specify the display when launching programs, for example:

```
xterm -display :0.1 &
```

launches an xterm in the second head.

If you are able, tell your window manager to ignore program supplied position hints. A multi-head aware wm will attempt to intelligently place the windows, otherwise, they tend to display in the 'centre' and straddle the heads.

Xinerama

The *xinerama* extension to XFree86 version 4.x allows a single, large desktop to be displayed across multiple heads. It is video driver independent, meaning that if your video card is supported, it will most likely work with *xinerama* (no guarantees, though). Assuming that you would normally use a 1024x768 screen size, *xinerama* in a dual head display would allow a screen of 2048x768 and display it across two monitors sitting side by side. Your mouse cursor will travel between displays and windows can be dragged across displays.

Ideally, that part of the desktop under the physical monitor edges would be resistant to windows expanding beyond them and dialogues would pop up centred in the physical monitor. We say "ideal' because the limiting factor to all this is the window manager. The placement of windows is controlled by the window manager, not X, so the wm needs to be *xinerama*-aware.

Hardware required

Linux has always been finicky about hardware and X is no different; multi-head even more so, and there are a couple of issues. A big issue for some people is that 3D support is not well-supported in multi-head mode or even available in most instances. This is a problem with the drivers, not the hardware, and until it becomes a priority for the development team (or you, and then you contribute you work to the community), greatly improved support is not likely to appear; a workaround is easy to implement (see below). Some video cards are not supported, and some are not suitable for use with multi-head; the reason having to do with how the video BIOS is initialised.

Common wisdom has it that Matrox cards are better for multi-head displays than others (although many others work) for the reason that some have been designed with multiple displays in mind, especially their G450 and G400 multi-head cards, and driver support is good. We found that it was necessary to update the video BIOS of an older Matrox Mystique card we purchased second-hand in order for it to behave properly. Additionally, this particular card had a place for a switch on the board to enable BIOS support for multi-head, only the switch was missing; an examination of the circuit and the use of a soldering iron solved our problem, but details on that kind of hardware hacking is beyond the scope of this article. Just be aware that your inability to get a particular video card working in multi-head may be the card itself. This is a common experience as we have observed that many of those attempting to enjoy the multi-head experience do so with video cards and monitors found in the junkbox, those having been discarded during previous »

The mysteries of XFree86 revealed.

Glossary

Head – a monitor.
Display – an environmental variable representing the name of the X server. Used synonymously, but incorrectly, with monitor. The displays are named using this convention:
localhost:0.0 – first display
localhost:0.1 – second display
and so on. The word 'localhost' is replaced with an IP address for remote displays.
Multi-head – using more than one monitor at the same time.
Dual head – the most common example of multi-head using two monitors.
xinerama – a special case of multi-head display where a single virtual desktop is created that spans all the displays.

Which video card?

There are several options for identifying the video cards in your computer. All work pretty much the same; our preference is XFree86 -scanpci. The goal is to discover the Bus ID so that this information can be used in the XF86Config-4 file.

Multi-head caveats

You must use XFree86 version 4.x – the newer, the better. Especially if you are experiencing problems with your display.

You distribution may not have the appropriate extensions compiled into the default XFree86-4.x binary. You may have to re-compile X.

Your video cards must be supported and should have the most current BIOS available if their BIOS is upgradeable.

You must run all displays at the same colour depth for *xinerama* (but not for independent displays; resolutions may all be different in any case).

Some older monitors or video cards may not support the resolutions, colour depths or scan rates you desire.

Due to driver limitations, a depth of 16bits is the best you can hope for with some cards.

Sadly, just because the hardware works in the Windows or Mac operating systems doesn't mean it will work in Linux.

Some window managers (and some older versions of current multi-head aware window managers) are not multi-head aware.

Even multi-head aware window managers have some variation as to how 'aware' they actually are.

Even with multi-head aware window managers, some programs are not multi-head or *xinerama* aware and will not behave as you think they ought to (Netscape being one notable example).

3D support won't usually work in *xinerama* mode and may not work in multi-head mode – it's the drivers, see.
Scared away yet? No? Good.

Dead space

If you attempt to enable *xinerama* using a 21" monitor and that 14" SVGA monitor it originally replaced, you may encounter "dead space" when you configure different screen resolutions for each of your monitors.

Let us assume that you configure a resolution of 1024x768 on the 17" monitor and 800x600 on the smaller one (because that's as large as it allows) and the monitors are sitting side by side. The virtual desktop *xinerama* creates is now 2048x768, but the smaller monitor doesn't cover all of its 1024x768 area; the "L" shaped space is the "dead space". A window manager may unwittingly open a window or display a dialogue within that space, effectively out of your reach. Different sized monitors and *xinerama* don't mix well without some annoyance.

« hardware upgrades. Monitors, to a lesser extent, may be culprits if they don't support the resolutions and scan rates you would like (or need) to use.

See 3D

As mentioned before, there is a workaround for using multi-head displays without letting all that costly 3D power go to waste. Fortunately, multiple XF86Config configuration files can be used with XFree86. By default, XFree86 will look for the default configuration files in /root, /etc/X11 and a few other places. Version 4 of XFree86 will look for a file named XF86Config-4 first and, if it can't find one, uses XF86Config. This can cause a problem if you have both versions of XFree86 installed on your system and aren't aware of this feature. Actually, XFree86 is so flexible, it allows you to use any config file located anywhere if you want to specify it with the **-xf86config** command followed by the full path and filename of the appropriate configuration file. This command can be passed to to XFree86 through startx by using the double dash command, so an example would look like:

```
startx --xf86config /etc/X11/XF86Config.dual head
```

Since both version 3 and version 4 can co-exist on the same system, you can run X in single head mode using the accelerated 3D graphics driver under XFree86-3.3.x to play games and then restart X with XFree86-4.x and your 2D multi-head configuration to actually get some work done.

Carpe Deim

Configuring multiple monitors is not easy, but not overly complex either when you possess an understanding of your hardware and the intricacies of X. The main stumbling block for ease of use is better driver support and better software support. These will come in time, especially if demand for this feature increases. We have found older 4Mb PCI video cards at reasonable prices and even the comparatively newer 16Mb Voodoo3 PCI cards are affordable (and even the Voodoo2 cards can be used for the second head). Of course, the dual-head Matrox cards are more expensive, but enjoy Linux driver support. Concurrently, 17" monitors are becoming more affordable and even the prices of LCD displays are coming within the grasp of the average person; 15" monitors are very affordable.

As long as you have safely saved your original working XF86Config file, you have little to lose by trying. Dive right in and expand your window on the world.

Multi-Head 12-Step Program

1 Make certain that you will be booting into a runlevel where X is not automatically started or running. This makes life simpler for a while if you normally start that way.

Just edit the line in /etc/inittab to read: **id:3:initdefault:**

or, you could use **linux 3** at the LILO prompt.

2 Backup and uniquely label any working XF86Config files you may already have, then delete all XF86Config files from the default location (usually /etc/X11) and shut down the computer.

3 Install all the cards you will be using in the AGP or PCI slots they will ultimately reside in. For those with AGP cards, the system BIOS usually has a selection to make the AGP card the primary video card (the one initialised first). Depending on the other cards you use, you may have to play with this setting.

4 Boot the system, observing which monitor/video card is initialised for the console display – that's your primary display. The other displays will remain unitialised and blank. If everything is blank, you have a card conflict (they are fighting over which will be the primary display. Remove cards one at a time until it works. From the command line, using you choice of **cat /proc/pci**, **lspci** or **XFree86 -scanpci**. The last one is recommended

you should be able to determine the Bus ID of each video card. Write down the Bus ID's of the cards.

5 Shut the computer down, remove all but the primary video card and restart the computer.

Using whatever tool you desire (we suggest 'XFree86 -

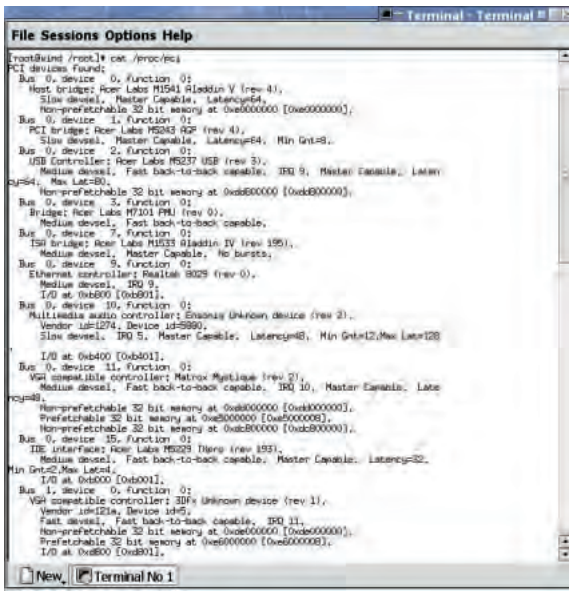

You can use 2D and 3D features by creating multiple XF86config files.

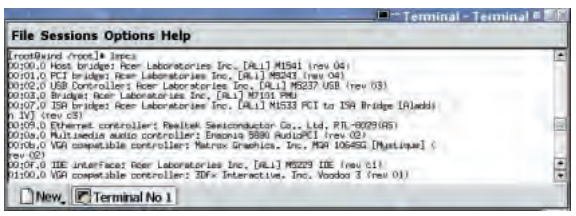

You'll need to configure each of your monitors.

Web Resources

XFree86 – http://www.xfree86.org
Xi Graphics – http://www.xig.com
Metro-X – http://www.metrolink.com
x2x – ftp://ftp.digital.com/pub/Digital/SRC/x2x/
VNC – http://www.uk.research.att.com/vnc/
x2vnc – http://www.hubbe.net/~hubbe/x2vnc.html
xmx – http://www.cs.brown.edu/software/xmx/
xmove – http://www.ensta.fr/internet/unix/misc/xmove.html
The Xinerama HOWTO – http://www.linuxdoc.org/HOWTO/Xinerama-HOWTO.html
XFree86 Documentation – http://www.xfree86.org/support.html
The XFree86 HOWTO – http://www.linuxdoc.org/HOWTO/XFree86-HOWTO/index.html
Remote X Applications HOWTO – http://www.linuxdoc.org/HOWTO/mini/Remote-X-Apps.html
Securing X Windows– http://ciac.llnl.gov/ciac/documents/ciac2316.html

Multi-monitor methods

Xnest

This tool allows you to run an x-server inside a window in an already running x-server. While mainly used for testing, this can be pretty interesting since you can launch xnest and start a complete desktop in it. Since we already have a :0 server running, let's start Xnest as the second server with: **Xnest :1** and then open an xterm in it with: **xterm -display :1** From that xterm, you may launch any app you desire or even open a new window manager and desktop. Essentially, this is what vnc is doing.

X2X

While not technically multi-head, *x2x* allows a computer running X to control, with its mouse and keyboard, another computer running X. To be useful, the monitor for the second computer should be on the desk next to your monitor. The source code is available from ftp://ftp.digital.com/pub/Digital/SRC/x2x/.

Assuming 'wind' is the current machine and 'marvin' is the target machine whose physical screen is to the left, the command line would be:

```
x2x -to marvin:0 -from wind:0 -west
```

X2VNC

If you happen to have both a Windows box and a separate Linux box, you can keep your desktop less cluttered and have control of both machines from the keyboard and mouse attached to your Linux box using *x2vnc*, a variant of VNC. The vncserver must be running on the Windows machine. You can obtain *x2vnc* as source only. Follow its **README** instructions to compile it and to launch it. Be aware that running most games on the Windows box will break the network connection to *x2vnc*. Use smbmount to make the Windows box's partitions available to your Linux box. The homepage for vnc is **http://www.uk.research.att.com/vnc/** and the homepage for *x2vnc* is **http://www.hubbe.net/~hubbe/x2vnc.html**.

The vnc server must be running on the Windows machine, and assuming 'wind' is the Linux machine and 'marvin' is the Windows machine whose physical screen is to the left, the command line would be:

```
x2vnc -west marvin:0
```

Xmove

Are you running multi-head, but not *xinerama*, and want to move a running program to another display? That can be done with *xmove*, a pseudo-server. The way it works is that it starts a server on :1 (your first server is on :0).Then, any client sent to :1 will display on :0, but can be moved to another networked computer. Not quite as easy as a **Send Window To** dropdown menu, but useful if you have the need. The homepage is at **http://www.ensta.fr/internet/unix/misc/xmove.html**; and the source is available from **ftp://ftp.cs.columbia.edu/pub/xmove**. After gunzipping and un-tarring the source, you need to use **xmkmf** and **make** to compile the binary. Take a careful look at the **README** files before beginning. After compiling, you will have two binaries, xmove and xmovectrl. The first starts the server, the latter moves the window to a different machine.

Start *xmove* from a console with: **xmove** and you now have a display :1 on your computer. Launch an xterm to it with: **xterm -display :1** and an xterm will appear on your desktop.

XMX

Developed at Brown University, *xmx* is 'an X protocol multiple-server' In essence, it shares an X display among a group, for example students in a classroom. The latest snapshot is from two years ago, but you could still find it useful in a classroom setting. The homepage is at **http://www.cs.brown.edu/software/xmx**. Binaries for Solaris, AIX and Linux are available as well as source. The package includes a selection of practical examples in the README files.

Xhost +

The easiest way to permit a computer to receive X clients over a network is to use *xhost +* to allow it. However, the connection is very insecure. A detailed description of more secure approaches can be found in the Remote X Apps mini-HOWTO at **http://www.linuxdoc.org/HOWTO/mini/Remote-X-Apps.html** as well a Securing X Windows at **http://ciac.llnl.gov/ciac/documents/ciac2316.html**. This site also includes a nice overview of X as well.

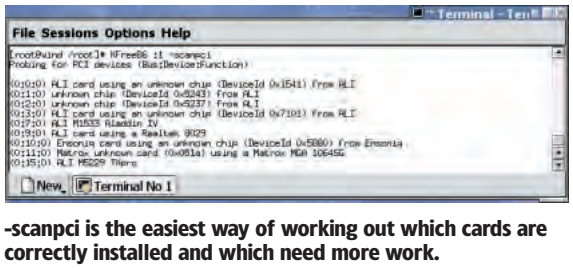

-scanpci is the easiest way of working out which cards are correctly installed and which need more work.

configure' for version 4 configuration), create a working XF86Config file (for both versions 3 and 4, if you have a need for it) for a single head display.

Since you are working as root, the test config file will be found in /root, your home directory, upon creation. It is a good idea to test the new config file with **XFree86 -xf86config /root/filename**. This starts X bare, without a window manager, so we can determine if X is configured properly without having to worry if the window manager is working properly. You'll see a grey screen with a large X cursor. After you get a working configuration, you can see how it works with the window manager by using: **startx --xf86config /root/XF86Config**.

Once it's working to your satisfaction, save this configuration file with a unique filename, then delete any XF86Config files remaining to start clean next time.

6 Repeat for each video card, one at a time (remember that AGP BIOS setting, if necessary!). The goal here is to confirm that each video card and monitor works and initialises at the resolution and colour depth you desire for the multi-head system (we suggest 1024x768 at 16bpp as a start – you can experiment later, after you get it working). Save each config file with a unique filename as before.

7 Now reinstall all the video cards where you had them before, reboot and run **XFree86 -configure** one last time. Move the config file to /etc/X11 and name it XF86Config-4. For now, make certain it's the only XF86Config file there.

8 Edit the file and ascertain that the Bus IDs are correct for each video card, that all cards and monitors have a unique section. Verify that the section describing the mouse is correct; we had difficulty getting the config program to recognise our USB mouse and had to cut and paste a working description from a working config file. Examine the section that describes the physical arrangement of the monitors (how they sit in front of you).

9 Run X naked with the command **XFree86**. Hopefully, X will start in the primary display and you will then see the secondary displays initialise and display the classic grey X screen. Your mouse pointer should move across all the heads.

If there is a problem, some help can be found by examining the error messages written to /root/.xsession-errors. If there's nothing there, check /var/log.

Compare the information in the config file with that in a working config file and adjust as necessary.

10 At this point, you should be able to start X with the default window manager using the **startx** command (the specifics as to configuring the startx process for a specific wm and a specific distro are beyond the scope of this article, thank heaven). You should now see multiple independent desktops, but may not see a desktop in each monitor if the wm is not multi-head aware.

11 If you want *xinerama*, test it with: **startx - +xinerama**

Remember to have all the colour depths the same. Screen resolutions can be different to accommodate different size monitors, just be aware of dead space.

12 You can *xinerama* the default by adding: **Option "xinerama" "yes"** to the ServerLayout section of XF86Config-4.

Which X?

XFree86 comes in two versions: version 3 and version 4. Some modern distributions install both, some only one. There are advantages to running version 3 if you need accelerated 3D support in some instances or have an older video card not supported in version 4. Version 4 is required to support a multi-head setup. Which do you have installed? Running XFree86 -version from a command line will disclose the version of X installed on your machine.

LOOK 'N' FEEL

Themeing WindowMaker

X is grey, but it doesn't have to be that way. Biagio Lucini shows you how to stamp your personality all over your Linux desktop using WindowMaker.

The setting of the background image is the first step towards the creation of a new theme.

Having a comfortable computing environment is fundamental to your productivity. And a major part of the working environment is the window manager (bearing in mind the old adage that you spend more time looking at your monitor than at your spouse). One of the best managers for Linux is *WindowMaker*, and through this tutorial we will help you to increase your level of comfort by creating your own theme. Granted, *WindowMaker* comes with plenty of themes, so in principle there is no need to spend time making your own. However, it is quite likely that none of them completely satisfy your tastes; perhaps you have a special picture you'd like to integrate into your workspace; maybe you want your screen to reflect your personality. Whatever your reasons, themeing not only provides a way to do this, but can also give you valuable insights into the way your desktop works.

Starting WindowMaker

We had a tutorial on *WindowMaker* in issue 12. In case you missed it, we will briefly restate you how to acquire and launch the application. First of all, make sure that you have the package installed with the usual procedure. For instance, if you are using a rpm-based distro, type **# rpm -qa | grep -i maker**. If you do not have *WindowMaker* installed, you'll probably find it on your distro CD, or download it from www.windowmaker.org. Once you have got it, run – as a user – the command **wmaker.inst**. The next step depends on how you are using X. If you start directly with X, you should add the command **wmaker** to your kdm, gdm or whatever login manager you use. If you start in Console mode, edit your .xinitrc file and replace the command you are using to start your present desktop environment with **wmaker**. Next time you start X, you will be running *WindowMaker*.

What is a theme?

Before starting to create our theme, we have to clarify what we are going to do. In general, a theme is the ensemble of all the components that determine the look 'n' feel of your desktop. What they actually involve depends on your choice of window manager. In the case of *WindowMaker*, we can identify the following principal components:

- the title bar, on top of each window
- the resize bar, at the bottom
- the text describing the windows
- the appearance of the menu (style, text etc.)
- the inconify and close buttons on the title bar
- the tiles, which are the common background of the icons
- the background image

Most of these components can be replaced; for instance, you can create your own tile set, or grab some cool tiles from the 'net (see the links page at www.windowmaker.org). We will assume that you have already all the components you need (You should at least have the default *WindowMaker* files to view and change) and the tools (*The GIMP* or *PhotoPaint* are ideal) to play around with their appearance.

Setting the background

First of all, we need to set the background image of our theme which will dictate the feel of the rest of it. Most of the common image formats are suitable for this purpose, just remember that huge, uncompressed photos of your dog will sap system resources worse than *StarOffice*. If you don't have any images to use, you can find a vast selection of them on the 'net covering every subject you can imagine. One of our favourite places for free desktop art is x.themes.org, but this is by no means the only place where you'll find good images.

Our new desktop will be based on LinuxSleep.jpg from x.themes.org. In order to set this image as background, copy it to ~/GNUstep/Library/WindowMaker/Backgrounds. Restart the window manager (from the menu, that can be accessed by right-

clicking on the background, select 'exit' and then 'restart'). Next select the image from the menu menu – Sessions –> Window Maker –> Appearance –> Background –> Images –> LinuxSleep.jpg.

Colour and style

Once you've set the background, the next step is to edit the colours of the various components (style). As a starting point, you can use any style (for instance, the style of the default theme is adequate) as a basis. The utility that helps in this operation is WPrefs. It should be present as a dockapp on your desktop and can be started by double-clicking the corresponding icon. If this is not the case, it can be invoked by the command **<path>/WPrefs** Just replace <path> with the path of the application. For instance, on our Mandrake box, it becomes /usr/X11R6/lib/GNUstep/Apps/WPrefs.app/WPrefs. This application can modify almost all the settings of *WindowMaker*, not just the style. Its interface is polished and intuitive.

Once the application has been started, it is convenient to enable (if it is not enabled) the 'Balloon Help' option. With this option on, mousing over an icon will cause a popup description of that icon's purpose to appear.

The first section we are going to look at is the one associated with the Balloon Help text 'Menu usability related options'. This is

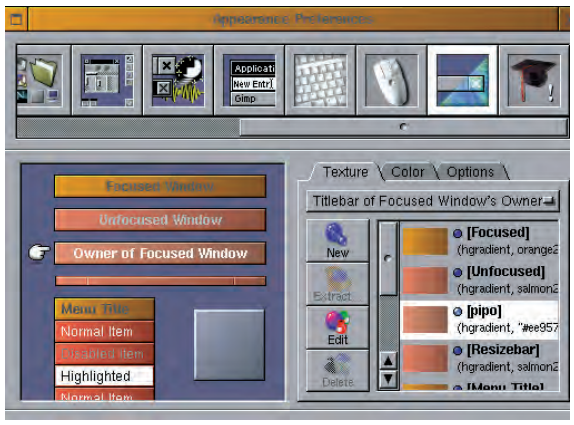

The Texture editor in WPrefs.

where you can set the relative alignment of submenus. Then we move to the section 'Animation speeds'. Among other options, we are interested in the 'Titlebar Style'. Simply choose between the two options.

The key point is the choice of the colours. To change the default, we move to the section 'Background texture configuration'. In this section, it is possible to configure all the components of a theme. The main preference window here is divided in two frames; in the right-hand side, it is possible to configure various parameters such as the texture and colour of window borders and scrollbars, while the left pane will show a preview of the customisable components. The procedure is straightforward: for the texture, select one component on the right-hand side by clicking on it, then select 'Edit'. You can choose between either solid or various types of gradient textures. The colour can be changed by clicking on the preview in the same window. An editor will appear that will offer you four methods of choosing colours: a colour disk with brightness bar; red-green-blue quantity editor; spectrum and predefined colours. Just choose the one you're most comfortable with. When you feel happy with the colour setting, click 'OK' and move to the next element. If the new setting does not show up, simply restart the window manager. In this way, you can modify each element of your theme, component by component.

Having chosen all the colours for textures of the windows and the menu and for the icons, it is time to move onto the text. Here things are even easier. In the frame 'Colour' on the right, select an item from the menu, then drag one of the predefined colours to the preview window and save. Again, restart *WindowMaker* if you would like to see the new settings at work. Finally, in the section 'Options' you can decide the style of the menu and the alignment of the titles. And that's that: your personal, unique theme is ready for action. You might even want to share it with the rest of the World (see Boxout, above).

Saving the theme

How can we save the present settings? *WindowMaker* has a dedicated and powerful command line utility called 'getstyle' for this purpose. To save the settings, 'getstyle' must be invoked with the **-p** flag. e.g **getstyle -p ~/GNUstep/Library/WindowMaker/Themes/LinuxSleep**. This will create the directory ~/GNUstep/Library/WindowMaker/Themes/LinuxSleep.themed (For further information on the getstyle command, type **man getstyle**). You should now be able to access your freshly saved theme by selecting: menu –> Session –> WindowMaker –> Appearance –> Themes –> LinuxSleep. LXF

The new theme at work.

Sharing and distributing your theme

If you're particularly proud of your creation, why not upload to the web, at wm.themes.org, and let the world share and appreciate your artistic genius. First of all, you need to package your theme in the standard format. Move to ~/GNUstep/Library/WindowMaker/Themes/ and then to your theme directory (LinuxSleep.themed in this article). Here you can place a file LinuxSleep.lsm in which you should give the following information:

- **name of the theme**
- **author and e-mail address**
- **description and copyright of the material used**
- **license of your theme**
- **other information you believe is needed**

Next, move to the parent directory and create a .tar.gz package: tar cvzf LinuxSleep.tar.gz LinuxSleep.themed. Then, connect to wm.themes.org and register.

To upload your theme, follow the instructions you find there (don't forget to prepare a decent screenshot and an apposite description of your work). If everything goes right, your theme will be available for download to all the *WindowMaker* users, who can benefit in this way of your work.

The ballad of VIC and SID

After last month's Spectrum overload, Simon Goodwin surveys Linux emulators for Commodore's successful 8-bit.

Commodore's 64 was the world's best-selling 8-bit computer. Introduced in 1982 to follow-up the million-selling VIC-20, it was a massive step forward in terms of graphics, sound and – most of all – memory capacity, with 64K of dynamic memory and 20K of ROM. That doesn't sound much – unless you're waiting for it to load it from an analogue audio cassette – but was a breakthrough at the time.

The 64 was first promoted as the Commodore Max, with an advertising campaign that trumpeted its elephantine memory, comparable with full-sized business machines. Indeed it was billed as "a sophisticated personal computer that's been specifically designed to run business applications."

The C64, as it became known, was clearly descended from earlier Commodore machines, back to the 1977-vintage KIM and PET. Yet it has survived into the new millennium, with production rights passing via Escom to Tulip.

Hardware

The C64 is built around two custom chips, VIC for video and SID for sound. Both were designed from the ground up in just six months in 1981. These very impressive chips proved most amenable to hacking. Programmers used them in many ways the original designers never intended. The simple system timing made it possible to synchronise processor, sound and up to nine overlaid images, gaining effects that were stunning at the time and still look pretty cool today.

Commodore's firmware was feeble by comparison. The 8K BASIC interpreter was bought outright from Microsoft for the PET, making the later machines cheap to produce but tough to program. The BASIC ignored the extra graphics and sound chips and could only access about half the 64K memory, the rest was hidden by the screen image, ROMs, and custom chip registers.

In the absence of multi-tasking – indeed, anything that we'd now know as an operating system – a typical C64 program is a lean, tricky, system-hogging tour de force. The only way to program the custom chips is to 'hit the metal', writing bytes directly to the array of 76 custom registers.

The Commodore 64's 6510 processor is closely related to the 6502 in the first Apple and Acorn machines. By the time the C64 came out Commodore owned the factory that made those processors, and they were so cheap that they were used as intelligent controllers in add-on peripherals, which meant great fun for hackers, rewriting the original slow serial interfaces.

The most common C64 expansion was the 1541 disk drive which, though slow and eccentric, was a big step up from cassette. C64 drives and printers use a proprietary interface, with a processor, rather than a peripheral controller, in each box. This idea came from the VIC-20 and earlier Atari computers. It ensured that Commodore made lots of money on add-ons, as well as computers. It also kept prices high and performance low, although third-party firms eventually muscled in with clone drives.

Software

Games, demos and music are the main attractions for most C64 emulators, but there are lots of PD utilities too, if you know where to look. Ironically, much as on the IBM PC, the limitations of the

Commodore 64 specifications

Processor: 8 bit MOS Technology 6510 at 1 MHz
Memory: 64K paged DRAM
1/2K static colour RAM
8K PET BASIC 2 ROM
8K KERNAL (system) ROM
4K Character display ROM
Sound: Mono, 3 channel 9 octave direct synthesis
Display: 320x200 pixels, 2 colours per 8x8 square
160x200 pixels, 4 colours per 8x8 square
40x25 character-mapped colour text
8 sprite overlays of up to 24x21 pixels
Maximum of 16 preset colours on any screen
Smooth hardware scrolling and panning
Hardware graphics collision detection

built-in firmware encouraged hackers to write much more capable operating systems. Like Linus, they felt they could hardly do a worse job than the manufacturers, and saw the potential for multi-tasking on the cheap and plentiful C64 hardware.

You're hardly going to run X, or even W, on a machine with a low resolution TV display and 64K RAM, but compilers and console applications are available for Unixoid systems like ACE, Asterix, Lunix and osT. Programmers still enjoy C64s, real or emulated, although the 40 by 25 character screen is a bit restrictive. You'll need the user manual, and preferably the Programmers Reference Guide as well, as emulator manuals assume you know how the C64 works.

C64 files

C64 files can be stored individually on Linux drives (with a small header containing system-specific information), on genuine C64 drives, or in special emulator files. D64 files are images of a whole C64 disk in one file.

The T64 format was also introduced by C64S, a PC shareware emulator. It contains a stream of files as they would be read from tape. PRGs and P00s are single-file formats derived from the native 'PRoGram' type, with the original name in a 26 byte prefix. PD programs like 'c1541' and '64copy' convert D64, T64 and P00 files for use with other emulators. 'wav-prg' converts between these and WAV files or analogue audio, for authentic cassette compatibility.

These are the key C64 commands needed to run programs from disk images:

LOAD "$",8 loads the directory from the first simulated floppy. Type **LIST** to see the directory. Wildcards are supported, so you can load the first file from the second drive with **LOAD "*",9,1**. The default behaviour is load into BASIC memory; the extra ',1' loads a file to the address at which is was saved from. **RUN** starts a BASIC program. Machine code is called by typing **SYS** followed by the decimal start address: **SYS 49152**.

You can use the arrow keys and backspace to edit and enter lines already on the screen to save retyping them. The key symbol locations may correspond to the layout on the original machine, not modern PCs, so check the documentation for your emulator – for instance the asterisk is where ']' appears on US PC keyboards. If some keys won't work, disable joystick emulation – it interferes with keyboard input on a real C64, and emulators faithfully mimic this

Outgrowing their calculator background, Commodore put a good-quality keyboard on the C64, though with only two arrows – emulators remove the need to press SHIFT to reverse arrow directions. C64s also has two nine-pin controller ports, normally expecting a digital joystick in port 1. Most emulators mimic these with the numeric pad or external hardware.

C64 device numbers

These are the numbers to be used in association with commands such as C64 BASIC LOAD, OPEN, PRINT etc.

1	**Cassette**
2	**Modem**
3	**Screen**
4-5	**Printer**
6-7	**Serial**
8-11	**Floppies**
12-31	**Expansion**

The emulators

We've found ten Commodore 64 emulators which can potentially run on Linux, including a couple written in Java. The best all-rounder is *VICE*, though *Frodo* can be more accurate in some circumstances, and EC64 has a number of unique features. Others include *ALE-C64*, *ComeBack64*, *CCS*, *X64* and *XCBM*.

Frodo

Frodo is well presented, well documented, and has plenty of features. Its rich supply of tweaked C64 ROM images means it's easy to get working. However, as it's a very accurate emulator, it's also resource hungry. *Frodo* aims to run graphics and demos more accurately than any previous C64 emulator, modelling every tick of the processor and graphics synchronously. Lesser

"Frodo aims to run graphics and demos more accurately than any previous C64 emulator."

emulators work a line or screen at a time, making them faster but inaccurate on programs that push the VIC chip hard. Thus, as the documentation puts it, '*Frodo* can display raster effects correctly that only result in a flickering mess with other emulators'.

Frodo comes in three flavours. The fastest uses line-based emulation, adequate for a good 80 per cent of games, but not programs that reprogram the VIC chip part-way through a line. *Frodo PC* is again line-based but supports hackish tricks like code ››

***Jetpac* meets *Defender* in Archer McClean's *Dropzone*.**

Animated titles from ace scrolling shooter *DropZone*.

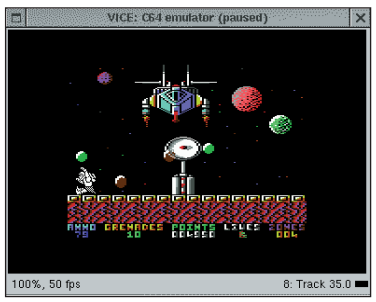

***Exolon* runs at full speed from an emulated 1541 disk.**

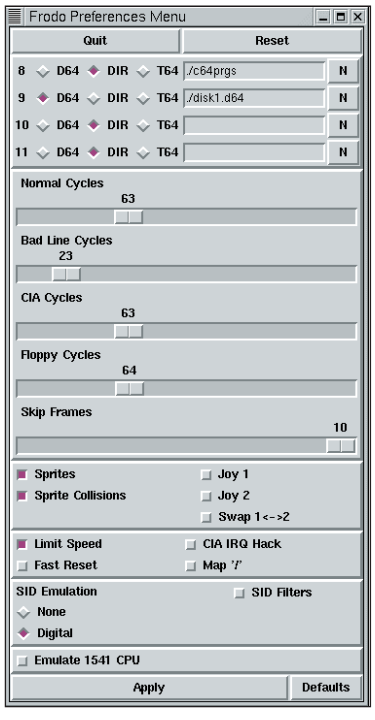

***Frodo's* TCL/TK front-end.**

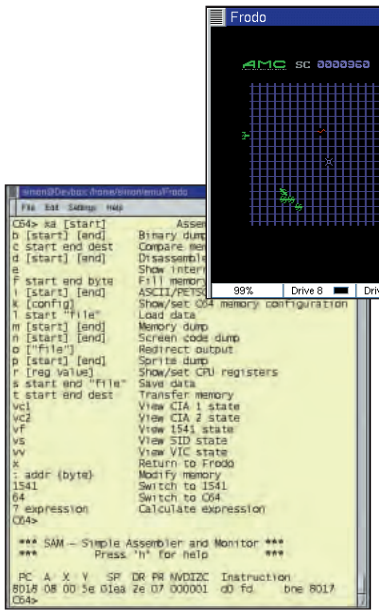

***Frodo's* 6510 processor debugging console.**

that runs inside custom-chip registers or relies on subtle aspects of the behaviour of the processor or interfaces.

Frodo SC is the slowest but most compatible. The letters stand for 'single-cycle' emulation indicating that the program switches from CPU to VIC emulation and back again two million times per second, locked in synchronicity with the original chips. Even this is not perfect, as the scant documentation note, but it's about the best you'll get, short of buying a real C64.

Not content with emulating VIC, SID, the 6510 and the CIA (Complex Interface Adapter) chip, *Frodo* supports the reprogramming of 1541 drives, emulating their processors and other hardware besides that in the C64 box. This is useful for disk software protection and fast loaders, although the need to run multiple CPU emulators slows things down noticeably.

Like the cycle-level VIC support, drive processor emulation is optional. Higher-level emulation of the 1541 loads most standard-format disk images almost instantly.

Each emulated drive can be set to access a disk or tape image or files in a Linux directory. *Frodo* supports X64 disk files as well as D64 images, and can read the native .LNX 'Lynx' tape archive format as well as T64 files. *Frodo* does not write tape images, and only writes to disk images when running precise emulation.

Frodo's machine-code monitor, SAM, can investigate and debug code for the 1541 as well as the C64. Drive activity and CPU speed is shown in the window border, with an optional TCL/Tk pane for adjustments to the configuration once it has started. Some of these have an immediate effect, while others are ignored till an 'apply' button is pressed. The initial configuration is read from a text file, and if unless you have the right version of TCL/Tk that's the only way to configure *Frodo*.

Our system had TCL/TK 8.0 but the GUI script refused to run, insisting on version 4.1. A comment in the TKGUI.tcl control file indicated that it should work with any version from 4.1 onwards, so we edited the version number to match the system and got the GUI up and running. The emulator and graphical controls still seemed semi-detached, and we found the surest way to make changes was to edit the text file and restart.

Mouse-controlled menus, like those in *VICE*, would give easier control, but we were able to select all of Frodo's options after some trial and error.

By default *Frodo* requires an 8-bit (256 colour) screen mode on Linux, though a badly documented patch is available for 16-bit systems. *Frodo* offers top-notch display, sound, hardware integration and speed control, but the user interface on Linux lets it down slightly.

VICE

VICE is an acronym for Versatile Commodore Emulator, and is the most comprehensive Commodore emulator for Linux. *VICE* can also run programs for other 8-bit Commodore computers, such as VIC-20 and C128, but the main focus is Commodore 64 emulation. With 300K of documentation in HTML and Info formats, and a similar amount of background text, it is comprehensively documented as well as implemented.

The only things missing are the vital ROM files, as the developers have not obtained permission from one of the incarnations of Commodore to include them with the emulator. They have done the next best thing by including a script file which attempts to download the files by FTP and install them for you.

VICE sets out for cycle accuracy, like *Frodo*, but uses a more pragmatic method which schedules 'events' only when it works out that pixel-synchronisation is needed, and runs less strictly the rest of the time, when it is not.

VICE is so packed with features – many of them unique – that we can only skim the surface here. It supports double-sized displays (in a window or full-screen), it has Linux joystick drivers, redirection for printers and up to four serial ports, and recognises an unparalleled range of disk images.

Besides D64 and its preferred X64 format, it can handle GCR, FM and MFM disk images that closely mimic the encoding of bits on other Commodore drives. Drive hardware emulation is an option on the first two disk units, 8 and 9. *VICE* can also read T64 tape images and P00 program files, and comes with C1541, a utility to convert between them.

You get a choice of two SID emulation engines – the newer one, Re-SID, is most accurate but takes a lot of processor time so you can switch between them, or disable sound, at will. *VICE* also benefits from an excellent UI and is easy to use and reconfigure on the fly.

The left button functions controls virtual drives, allowing disk and tape images to be swapped in and out. The right button adjusts other aspects of the configuration – speed, sound, graphics, and peripherals in general.

There's no need to compile *VICE* or install it manually, as it is available as a package, but the latest source-only release 1.6, which adds support for 'arTs' and 'esd' Unix

Sounds like SID

The custom chips are the real test of C64 emulation, and it takes a lot of processor power to impersonate them accurately. The designer of the C64 sound chip, Bob Yannes, went on to found the synthesiser firm Ensoniq, whose chips power PCI sound cards.

SID is a synthesiser, rather than a sample-player, though it can deliver grainy sampled sounds. It revolves around three phase accumulating oscillators. Each can play a distinct note in various timbres, with envelope shaping, filtering and ring modulation for variety. SID's 16bit frequency controls made it the first micro sound chip that could play in tune. This blew away the chips used in other early polyphonic home computers, which lost accuracy as the pitch went higher.

SID's distinctive sound stems from preset digital waveforms and hybrid analogue circuits. It was far more expressive than the square waves which were all the earlier chips could manage, though quality varies a lot between individual SID chips.

Good emulation is processor intensive, so many programs get by with basic emulation. The more demanding emulators make a big difference to the best modules and game music.

You don't have to run a full-blown emulator to listen to SID tunes. There are Linux libraries just for this, with a choice of console and graphical controls. SID-playing plug-ins also exist for Netscape and XMMS. The Sidplay download page lists a dozen variations on this theme.

sound abstractions. The other changes between version 1.0 and 1.6 are fairly subtle, but encouraging in a world where few programs ever seem to reach version 1.0.

Lunix is a credible multi-tasking OS for the Commodore 64.

There are other Linux C64 emulators worthy of consideration, especially if you'd prefer something relatively svelte, but *VICE* is the one which they all have to measure up against.

CB64

The CB in the name stands for ComeBack, and *CB64* is a port of the ComeBack64 emulator originally written for Wintel systems. It uses Marat Fayzullin's LibUnix portable emulation library, and runs on FreeBSD as well as Linux. It should also be portable to non-x86 processors.

ComeBack64 requires X to be running in 256-colour mode. Even then it insists on rendering graphics in greyscale on a colour X display. This seems to be a peculiarity of the Unix port – the new Java version, currently at alpha4, displayed colour correctly in Netscape 4.71.

The keyboard and user-interface of the Java version of CB64 need more work, but ironically it looks as if that might be the best version for Linux, unless the C port gets a swift overhaul.

Smart Alec

ALE-C64, also known as *Alec*, is a free emulator written in C for Linux, MSDOS and various Unices. The Linux version can run in an X window or full-screen via SVGAlib.

You get a dozen versions of *Alec*, with and without overscan display of the screen border, and with three flavours of audio: direct to the PC beeper, wave synthesis through a soundcard with OPL3, generally included to play MIDI tunes, or samples through '/dev/dsp'. That's the more capable but most CPU-intensive option, even though output is only 8-bit mono at half CD rate.

Sound is not *Alec's* strongest point – none of the emulators support the filters, ring modulation or exact synchronisation expected by the coolest C64 audio. The OPL and beeper output requires you to run *Alec* as root, so it can bash the metal of the host system. Results are better than nothing, but not on a par with *Frodo* or *VICE*.

The keyboard layout mimics that of the original C64, easy enough to use if you've had a real C64 and don't look at the keytops. Failing that, there's a diagram in the documentation. There are no menus but the startup and control key options are quite comprehensive and clearly explained.

Joystick emulation uses the keypad for either C64 stick. *Alec* does not support real joysticks, or disk images – it emulates the C64 but not the 1541 drive. Files are held individually on the host file-system, with a text file to associate local names, paths and loading commands. *Alec* can handle P00 program files and non-random data files, but not the commonplace collection formats, T00 tapes or D64 disks.

Alec requires a single file with the three Commodore ROMs in it. If these are not included in the archive it's easy enough to make those from the files supplied with *Frodo*, like this:

```
cat "Basic ROM" "Kernal ROM" Char ROM" > romimage.c64/.
```

ALE-C64 supports very fast and slow machines by limiting speed or skipping display frames. It generates images a pixel line at a time, which is a good compromise between speed and compatibility for most C64 applications. Generally, *Alec* emulation is pragmatic. It has small gaps but they're understandable ones that do not stop it being useable. It's greatest weakness is lack of direct support for common emulator file formats, but you can work around that with utilities like C1541.

EC64

EC64 is potentially the fastest emulator reviewed here, but also the most machine-specific. It was originally distributed as assembly-language source for NASM 0.98, the Netwide assembler, but the new version 0.17 is available as a binary that suits most IBM PC-based Linux systems. The hand-coded 6510 emulation core marks it out from the rest, but stops it running on Alpha, 68K, PPC or other non-Intel Linux flavours.

EC64 describes itself as experimental, yet it's shaping to be the dark horse of Linux C64 emulators. It already runs classic Reflex demos impressively, though disk loading is currently rather slow. Sound, sprites and joystick emulation recently arrived, and the documentation, though sparse, is adequate and improving.

Unless you have a framebuffer device *EC64* takes over your VGA hardware, so application switching is problematic until you quit the emulator, though it can grab screenshots. If all you really want is a C64, try the boot image, which puts a Linux system, the emulator and three demo disks on one 1.44 Mb floppy; no hard disk required!

EC64 is controlled with function-keys and a simple shell to change disks and adjust parameters, and a potentially powerful script language that can issue strings of BASIC commands as well as its own instructions. *EC64* is an elegant, expressive hack in the spirit of C64 programs – its small, at only 100k, and advances by leaps and bounds. The unique approach, scripting and floppy boot capability mark it out as one to watch.

Web resources

Old machines never die, they just retire into cyberspace.

C64 resources http://www.cs.cmu.edu/~dsladic/vice/vice.html
Home of VICE http://www.Uni-Mainz.DE/~bauec002/
Home of Frodo http://mars.wiwi.uni-halle.de/ec64/
Home of EC64 http://www.fatal-design.com/ccs64/
Home of CCS64 http://www.pure.de/thomas/
J64 – a C64 in Java http://c64.rulez.org/chromance/bbs/ftp.cgi
Gangsta's Paradise C64 Linux compilation http://www.primenet.com/~tbies/adocs.html
Tom's Commodore Emulation page http://www.npsnet.com/danf/cbm/
Dan's 8-bit CBM Information http://pulsar.c64scene.org/emulators/c64/
C64 emulator compilation http://www.heilbronn.netsurf.de/~dallmann/lunix/
Lunix – a Linux-like C64 OS http://www.nada.kth.se/~d93-alo/c64/links.html
Commodore 64 Nostalgia http://www.geocities.com/SiliconValley/Lakes/5147/sidplay

And that's not all...

X64 **is Jouko Valta's original Unix Commodore 64 emulator. He moved on to the VICE team, but X64 lives on as the command used to start the Commodore 64 part of that multi-emulation project. *X64* is still available for download, but won't win any prizes for features or compatibility.**

CCS64 **is a shareware emulator for Windows. An early version of this was ported to Linux in 1997, but the only surviving archive of this, at ftp.funet.fi, seems to be corrupt. The Windows version, now at release 2.0, is highly rated, but Unix compatibility seems to have drifted off the author's agenda.**

XCBM **is a very basic emulator which runs in a terminal window. It makes precious little use of X, and the appeal is probably confined to people who want to run CBM BASIC or text applications in a remote terminal. *XCBM* works, but the lack of sound and graphics make it more a retro curiosity than a mainstream C64 emulator.**

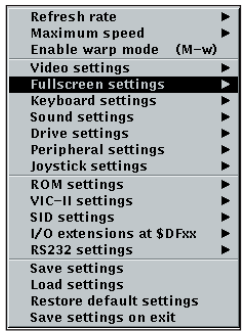

The left mouse button summons *VICE's* file handling menu.

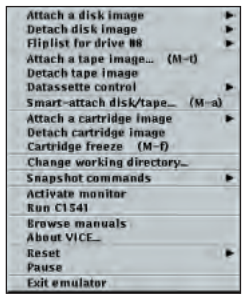

VICE **settings are accessed via the right mouse button.**

IMAGE ACQUISITION WITH SANE

Scanning made Easy

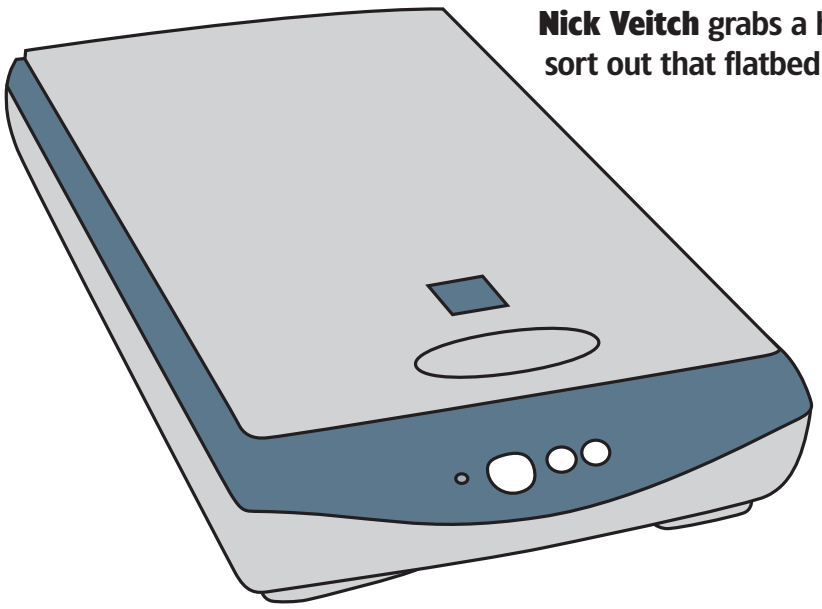

Nick Veitch grabs a handful of family snaps show you how to sort out that flatbed – once and for all.

Setting up a scanner seems to be a tricky issue for a number of people. Like most hardware in this Wintel biased world, little thought seems to be given to people who don't run Windows.

There are similar problems with printers, cameras, modems and other esoteric devices, but scanners conform to a standard, called TWAIN, which should make supporting them under Linux as simple as supporting TWAIN under Linux. Don't they?

Unfortunately not. The TWAIN interface system, while it does a good job of making the process simple under Windows, doesn't quite fit in to the Linux way of thinking. Mainly this is because the user interface is connected intrinsically with the driver. You may have noticed if you have used such software under Windows that the GUI on the TWAIN software varies depending on which company manufactured it. This simply isn't the Linux way, and also leads you into problems – if you wanted some scanner software that supported the latest features of *KDE 2.1*, for example, you'd also have to create a completely different one for those who didn't run *KDE*. Clearly, it's hard enough to get support in terms of one driver, never mind having half a dozen versions.

The Linux scanning solution, SANE (Scanner Access Now Easy) solves these problems by splitting the software into two parts – the backends, which are hardware specific, and the frontends, which can support any of the supported drivers. This is more flexible as terminal driven front-ends and network transparent interfaces can be constructed.

Important Note!

Parallel port scanners cannot be used unless you are running as root, or the scanner software is given root privileges. This is because the drivers actually have to talk to the parallel port directly through the kernel.

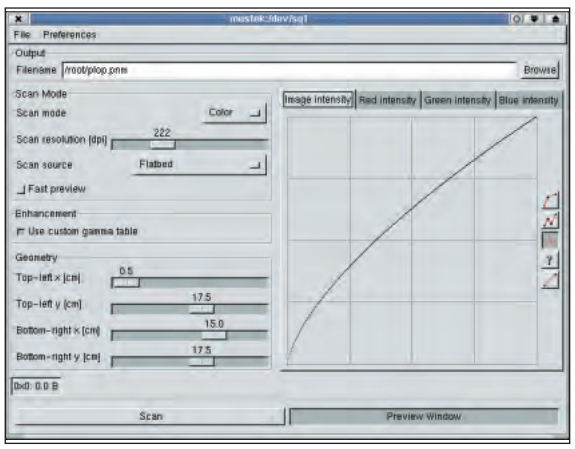

Xscanimage also supports gamma correction on seperate RGB channels for fine tuning.

Get SANE

In order to get SANE running on your system, the first thing you really need to do is get the SANE backend package. This is included on the CD this issue, but you may find that you have it, as it comes as part of popular distributions such as SuSE, Mandrake and Red Hat. To check if your distro has installed it on an rpm based system, you can type:

```
rpm -qa | grep sane
```

which should come back with something like 'sane-1.04-3'.

If not, you can install it from the packages on the CD or, to make sure you have the latest version, check out the main SANE ftp site at ftp.mostang.com/pub/sane/

If you need to install from source, follow the usual steps for compiling:

```
./configure
make
make install
```

You may also want to build the SANE-frontends components at the same time.

Check it out

Before you start to use sane, you'll want to test that the installation is working, and that it can spot your scanner. Fortunately, both of these can be done at the same time. Enter:

```
scanimage —list-devices
```

in a terminal window. This asks the SANE software to check what SANE supported devices it can find attached to your machine. You should get a response similar to this:

```
device 'mustek:/dev/sg1' is a Mustek ScanExpress 6000SP
flatbed scanner
```

which lets you know that your scanner has been found, and also whether the software has identified it properly.

If no devices are listed, there are a few likely reasons – either your scanner is not currently supported (check the full list at www.mostang.com/sane/) or the device itself is not being found. For SCSI devices check the /proc/scsi/scsi file to see if your SCSI device is being picked up. If you are using the SCSI interface that came with your scanner, this may well be the problem – some of them are built using incomplete SCSI capabilities, and may not be supported. Try to find out which SCSI chips the card is based on and check that it's supported by the kernel.

To make sure that the backend supports your scanner

Is my scanner supported?

There is a huge list of scanners currently supported by SANE. For the full list, you should check out the SANE website at www.mostang.com/sane, which is regularly updated. Just because your scanner isn't listed, doesn't mean it won't work, as many scanners are rebadged versions of main brands, or are functionally similar to existing scanners, even if the model numbers are different. Common scanners which are supported are:

EPSON
GT500, GT6000, GT5500, GT6500, GT7000, Perfection610, Perfection1200 (SCSI and USB)

Hewlett Packard
HPScanjet series: IIc, IIp, IIcx, 3c, 3p, 4c, 4p, 4100C, 5p, 6100C, 6200C, 6300C, and some members of the OfficeJet series.

Microtek
Scanmaker series: E2, E3, E6, 35, 45t, II, III, IISP, V300, V600, X6,330, 630

Mustek
Paragon series: 600S, 600 II CD, 600 II SP, 800 S, 800 II SP, 6000 SP, 8000 SP, 1200 SP, ScanExpress series: 600 EP, 6000SP, 1200SP, 12000SP

Plustek
OpticPro series 4800P, 4830P, 600P, 6000P, 9600P, 9630P, 12000P

AGFA Snapscan
300, 310, 600, 1236s

UMAX
Astra series: 600S, 610S, 1200S, 1220S, 2100S, 2200SU, 2400S

properly, you can activate the test mode with the command:

```
scanimage -T
```

Up front

The simplest front end for SANE is the *scanimage* tool you have already been using. This command line tool can also be used to scan directly. Enter the command

```
scanimage —device DEVICE_NAME —help
```

where the DEVICE_NAME is the full device name reported by SANE, e.g. Mustek:/dev/sg1 from the previous example.

This will give the the command line options available for your scanner, which will probably include processing defaults such as:
—brightness -100...100%
—contrast -100...100%
and so on. These vary depending on what type of scanner you own, and should include options to address specific features such as transparency adaptors as well. A typical scan might be performed with the command:

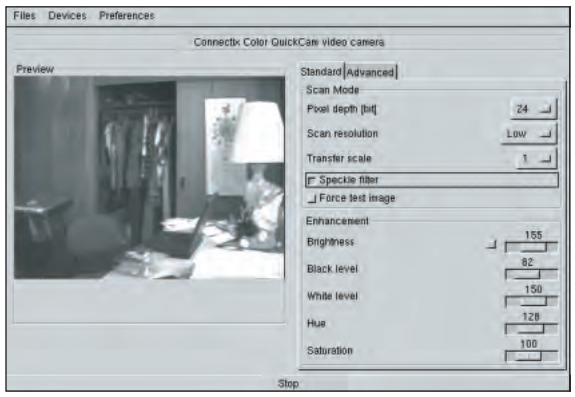

SANE isn't limited to scanners - this is a front-end for SANE's Quickcam driver.

```
scanimage —device mustek:/dev/sg1 —brightness 50 > output.ppm
```

Although you can use the command line for scans, obviously it's much easier to use a GUI driver program. The default frontend for SANE is the *xscanimage* program distributed in the SANE-frontends package. This supports the basic features you would expect (preview mode, dpi specification etc) and also advanced tools such as gamma correction.

You can run *xscanimage* directly from the command line, but you may prefer to run it as a plugin through *GIMP*. If you already have it installed, you can create a link to the plugins directory. First of all, discover where the binary for *xscanimage* has been placed. Then create a symbolic link to a file in the users .gimp folder. For example;

```
ln -s /usr/bin/xscanimage ~/.gimp-1.1/plug-ins
```

Now if you run *GIMP*, you should have an 'acquire image' item under the Xtns menu, which will launch *xscanimage*.

Other frontends are available, most notably *XSane*, which offers other features such as photocopy and fax options.

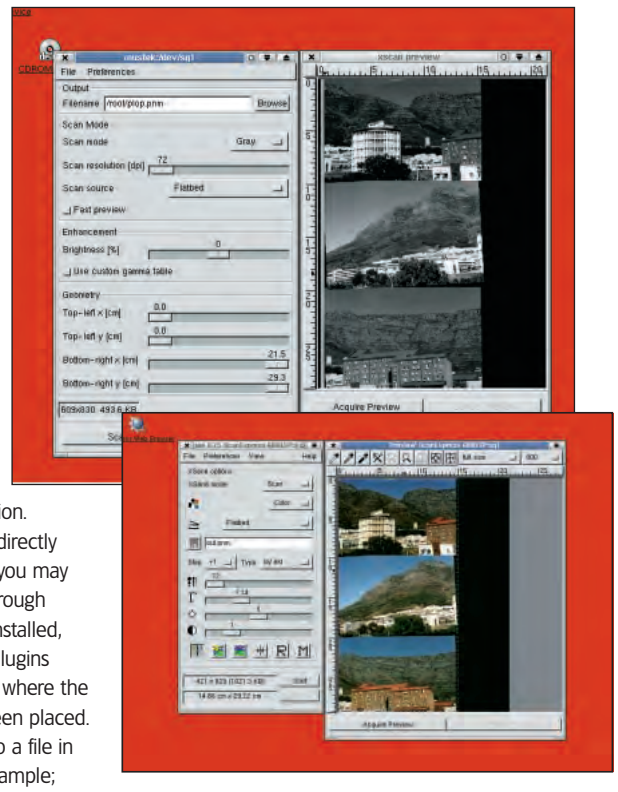

xscanimage (top) and XSane are just two of the frontend scanning options available.

Future enhancements

Of course new drivers are being written all the time, but SANE drivers are no longer being limited to scanners. As with its TWAIN counterpart on other operating systems, SANE drivers are being written for other image capture devices.

If you run into problems, or you want some more information on SANE, check out the quite comprehensive FAQ at http://www.xs4all.nl/~ljm/SANE-faq.html. LXF

Types of scanners

Good support depends on your scanner's interface.

Broadly speaking there are three types of flatbed scanners, which fall into categories depending on the interfaces used to connect to them to your PC – SCSI, Parallel and USB.

Support for SCSI scanners is good, mainly because the scanners have to comply with the command set used by SCSI devices to communicate. Unfortunately, although the commands themselves are universal, the way each scanner manufacturer decides to implement those commands may vary, so not all SCSI scanners are the same. Also, some scanner makers have added extensions to the command set. This is why some scanners are only partially supported in SANE (e.g. they may only work in greyscale mode).

USB scanners similarly have some commonality, but as USB support under Linux has been a long time in coming, work has only really begun on supporting these. Thankfully, because of their co-operation with developers, some companies have managed to make great efforts to make sure their USB scanners are supported. Practically all of Epson's USB scanners are supported by SANE.

Parallel port scanners are the hardest to create drivers for, as there is no common protocol or command set for accessing such devices. Although many manufacturers tend to use a similar way of driving all their parallel port scanners, there are a few exceptions, and the methods used vary widely between manufacturers. Much headway has been made into creating drivers for these scanners, but if you have an old or unusual parallel port scanner, you may be out of luck.

What on Earth is... SOAP?

SOAP is no longer just a product for removing grime. The Simple Object Access Protocol could herald a whole new 'netcentric way of working.

So, I take it this sort SOAP has nothing to do with my morning shower ritual then?

No, it's far more exciting than that! The latest buzz phrase to come out of the IT industry is Web Services. Briefly, this is the ability to access services such as remote disks, printers, objects (as in OO) and applications and so on over the Internet. Ignoring the security issues associated with this, the underlying technology is very interesting, but not just from an Internet point of view. There is no reason why these services, and the underlying protocols, data encoding etc cannot be used internally within an organisation.

The biggest player in this brave new world of connectivity is our old friend Microsoft, with their .NET strategy. One of .NET's aims is to provide the ability to access applications, data and so on over the Internet, obviously there will be a small charge for this (not forgetting security nightmares, data ownership issues etc). However, out of the .NET platform, an interesting method to remotely access objects in a fairly simple manner has developed. It's called SOAP.

SOAP allows you to create a distributed applications, and use standard Internet technologies to communicate with the separate objects that make up an application. This does not need to be limited to just object orientated languages such as C++ or Java, but can be used with any programming language that has the ability to communicate using SOAP. Currently there are 41 SOAP implementations available including those for PERL, Python, and PHP to name a few.

Does this mean the end to platform specific applications then, or just another set of unstandard standards to ignore?

The Simple Object Access Protocol (SOAP) is designed to be used to invoke services, components or objects on a remote server. SOAP's main objective is to provide the means to create distributed applications that, in turn, can use services that are themselves distributed. SOAP can be likened to COBRA or DCOM, however it has been designed to be more Internet friendly than these platforms, easier to use and possibly less restrictive. SOAP has been designed to use HTTP as the data transport protocol and XML for data encoding. Since both HTTP and XML are open standards, this should ensure that SOAP is not restricted to Microsoft sanctioned platforms.

A SOAP connection model consists of a SOAP client sending a request to a SOAP server, which in turn forwards the request to the required object for execution. Assuming that a response is required, the SOAP server will relay this response back to the SOAP client. SOAP can run in either request/response mode or just execution mode (i.e. no reply is expected back from the remote object).

When the SOAP client sends a request to a SOAP server, the XML communication contains the name of the remote object and the method that the client application requires to be invoked upon that object. The required method is contained within the body of the XML document, whilst the object is defined in a HTTP header called SOAPAction.

Actually, the name 'SOAP server' is a little misleading as the server does not actually perform any real work, apart from relaying the received request on to the object in question and potentially returning any resultant data. For this reason you will often see the SOAP server referred to as a SOAP proxy because of this relay function.

SOAP originally grew from a similar project, called XML-RPC. This project was designed to allow Remote Procedure Calling (RPC) communication to be performed over a network, again using HTTP as the transport and XML for encoding. This project is widely deployed and may be worth looking into if all you wish do is perform RPC over the Internet or wide area network where standard RPC mechanisms can not be used.

I've used COBRA or DCOM in the past. What does SOAP mean for these technologies?

The biggest issue with COBRA is that it is not always easy for one manufacturer's COBRA implementation to communicate with another manufacturer's version. This is because of the way each manufacturer implements the COBRA standard. Therefore, although on paper you should be able be run multiple versions of COBRA and the associated ORBs, in practice this is not always the case. This means a company normally has to standardise on one COBRA platform to ensure a successful solution. SOAP does not directly address this issue, but instead provides another mechanism for a COBRA application to access a remote object.

With DCOM, the biggest problem is that this is tied into the various Microsoft platforms. There have

been DCOM environments developed for other platforms, but these have not taken off, mostly due the fact the DCOM is a Microsoft platform and not an open standard, which COBRA can claim to be.

In addition to these issues, both COBRA and DCOM have not proven to be very Internet friendly, and therefore some other mechanism is required to access services on a remote server located on the Internet. This is not to say that SOAP is a replacement for either COBRA or DCOM, but instead provides a method for these platforms to access remote objects over the Internet, or wide area network (WAN).

Aren't there more efficient ways than HTTP to shovel stuff around a network?

There are several reasons why HTTP is the default protocol, although other protocols can be used if required. The most critical reason is based around firewalls. As SOAP is intended to provide a platform to access remote resources over the Internet, a protocol was needed that required no configuration of firewalls. As HTTP is usually allowed to pass through a firewall, using HTTP as the data transport allows SOAP to be used without any firewall changes. In addition, as HTTP is a W3C standard, interoperability between differing operating systems is easier to achieve.

Isn't it likely that Microsoft will use the old embrace and extend method to make sure only Windows users can make the most of it? It seems unlikely that Bill Gates has gone all open source on us.

SOAP's main proponents are IBM and, indeed, Microsoft, both of which performed the initial development work on the protocol. It is also true that SOAP is forming a strategic part of Microsoft's overall .NET strategy, but the SOAP standard is not owned by Microsoft, or any individual company. It's now owned by World Wide Web Consortium (W3C), which is the body responsible for recommending Internet related technologies. The W3C is responsible for the other Internet technologies such as HTML, XHTML and XML to name a few. This should help to ensure that no single vendor can change the SOAP standard to their advantage.

It is strange to see Microsoft activity involved in an open standard, however if you look at the future strategy that Microsoft has set itself, with the .NET platform, embracing an open standard makes sense for them. To ensure the success of .NET Microsoft need to ensure that a large number of platforms can access any .NET application. Whether this works in the real world remains to be seen, but the business logic makes sense.

That's the theory, but what if I actually want to do something useful, practical or entertainment?

As mentioned, SOAP use XML to encode the data content, which in then sent – via HTTP – to the remote server for execution. The SOAP specification defines this XML data structure and how SOAP in sent over HTTP, but these are beyond the scope of this article. As a quick overview, the following outlines a possible SOAP request to a remote SOAP server:

```
POST /LatestNews HTTP/1.1
Host: www.linuxformat.com
Content-Type: text/xml; charset="utf-8"
Content-Length: x
SOAPAction: "http://www.linuxformat.com/new-news.xsd"

<SOAP-ENV:Envelope
        xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
        SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
        <SOAP-ENV:Body>
                <m:LatestNewsHeadline xmlns:m="http://www.linuxformat.com/new-news.xsd">
                <search>Debian</search>
                </m:LatestNewsHeadline>
        </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

Obviously this is a very simple example, however it shows a basic SOAP request. In this example, a request would be sent to http://www.linuxformat.com/new-news asking for the latest news on Debian. As mentioned previously, the SOAPAction header defines the remote object that the request is direct to. The SOAP-ENV:Envelope is required and defines the SOAP headers and body to identify the method to be invoked on the object, with any additional parameters.

When the SOAP server receives this request, it parses out the contents and then invokes the object on the local system. Once a response is received this is then passed back to the client. The SOAP response to this example request might look like the following, returning the element LatestNewsHeadlineReponse which contains the latest news for Debian:

```
HTTP/1.1 200 OK
Content-Type: text/xml; charset="utf-8"
Content-Length: x

<SOAP-ENV:Envelope
        xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
        SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"/>
        <SOAP-ENV:Body>

<m:LatestNewsHeadlineResponse xmlns:m="http://www.linuxformat.com/new-news.xsd">
                <HeadLine>Debian releases 2.2r3</HeadLine>

</m:LatestNewsHeadlineResponse>
        </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

Sounds great: a safe, stable way of working online, the culmination of all that 'the network is the computer' stuff that Sun were bleating on about ages ago... Hang on, it is stable, isn't it?

SOAP is still considered by some to be a work in progress, with some issues around XML Schemas. The actual difference between SOAP 1.0 and the current version, which is 1.1, is very limited and appears to have more to do with IBM joining the party than any major changes. However, manufactures are now implementing SOAP version 1.1 support into the product line with Microsoft heading the charge. The next version, 1.2, is currently being worked upon but is not expected to be approved for awhile.

I've searched the 'net for SOAP and keep getting directed to sites extolling the benefits of regular washing. If it's the Next Big Thing, why is no-one talking about it?

Surprisingly there is not all that much information about SOAP on the Internet. Currently, the best place to start is at Microsoft web site (sorry about that), but O'Reilly's xml.org site is also a good place to search for documentation and has links to several other sites. The easiest route to start to play with SOAP is to look at the PERL module SOAP::Lite or Python's soap.py as this will give you a grounding in SOAP and allow you to review the practical uses of this technology. LXF

Find out more...

There are a number of places where you'll find more information on the Simple Object Access Protocol, you just need to know where to look.

http://www.w3.org/TR/SOAP/
Here you'll find all the latest news about SOAP, and as much information about its uses as you could possibly need.
http://www.soapware.org/
This self-styled 'leading directory for SOAP 1.1 developers' contains extensive information (and links) to all of the current implementations.
http://www.microsoft.com/net
Discover Microsoft's plans for the future.

LINUXformat Developer

All the latest tools, info and advice for Linux hackers

THIS MONTH...

NEWS

New GUI toolkit

QT 3.0 gets closer...

The popular cross platform C++ toolkit Qt is building up to its 3.0 release. Qt is the basis of the *KDE* desktop environment, and the new release includes a host of new features and facilities. Improvements include a substantially improved Qt Designer (including editor so you can code within it), platform independent database support (Qt is cross platform already), multiple monitor support, new component model, a new regular expression engine, accessibility support for visually impaired users, expanded controls and a documentation browsing tool called Qt Assistant. This improved release should pass on its benefits to the *KDE* project and we can expect to see *KDE* take on board these extended facilities in a future release.

Although Qt 3.0 has been in development for a while, the connection between the *KDE* team and the Qt team has been quite intimate in details. Code contributed by *KDE* developers has always been available to improve Qt, and in return the improvements to Qt improve *KDE*. An example recently was when the *Koffice* developer David Faure backported a feature in Qt for use in *KWord*. This connection between the commercial and free software communities has worked well.
http://www.trolltech.com/

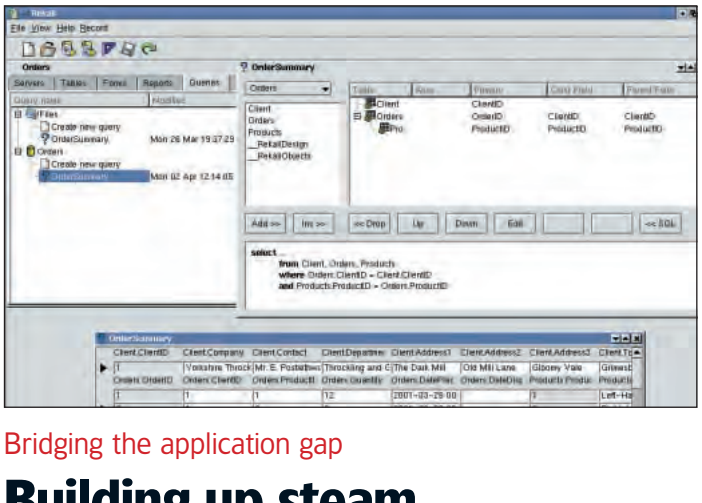

***Rekall*, theKompany's Access-style database, is one of a number of imminent releases.**

Bridging the application gap

Building up steam

theKompany.com, a publisher providing applications and development tools that natively run on Linux, have been working furiously on a new range of products aimed at desktop users and Linux developers.

Planned releases will be attempting to bridge some of the gaps in the Linux software market and include *Aethera*, a personal information manager and groupware suite; *BlackAdder*, an integrated development environment (IDE) for python-based development; *KDE Studio Gold*; a professional version of the free IDE for *KDE* which extends the software with value added features such as extended documentation, code editing and development facilities; *Kapital*, a personal financial manager and Powerplant, their development environment for popular Linux distributions.

The software is mainly written as native *KDE* applications, so the products integrate well into the *KDE* environment. theKompany have also worked hard to give back to the Linux and *KDE* community by contributing three applications to the *KOffice* suite. *Kivio*; a professional *Visio*-style flowcharting application; *Kugar*, an XML reporting tool; and *reKall*, a long awaited *Access*-style database development package. theKompany also have a range of commercial stencils available for purchase which provide a wide range of shapes for use in *Kivio*.

This extension of free software is a novel business plan that goes beyond the typical support and consultancy business plans of other Linux companies, although they do offer consultancy, custom software production and porting services.

So far, managers and coders at theKompany have worked well to identify the areas where Linux falls down at the desktop, and their solutions – as they head for final releases – should be interesting to see. As development of the various products continues, it will be well worth keeping an eye on them if you will be using Linux for development or desktop productivity.
http://www.thekompany.com/

Comment

Crash protection

SGI releases XFS 1.0 journalling filesystem

SGI have released their XFS journalling filesystem software for Linux. It has taken two years for SGI to get XFS up to a production release, and its release will mean lengthy bootups after a crash (when fsck scans and fixes your disks) should be reduced.

XFS is not the only journalling filesystem available for Linux, with IBM's JFS and ReiserFS already gaining converts. SGI has also released *FailSafe* alongside XFS. *FailSafe* is a clustering tool capable of servicing 16 nodes, with each node capable of sharing storage in the event of another nodes failure.

Slow bootups on Linux have been one of the complaints from new users, and the advent of GPL'd publicly available free journalling filesystems means these slow bootups should be a thing of the past. Although the SGI and IBM filesystems will reduce bootup times, the new ext filesystem is in development (ext2 is the current filesystem type that most Linux machines use). Ext3 has been in development by Dr Stephen Tweedie, a kernel hacker who has been contributing to the kernels filesystem and virtual memory code since the early days.

SGI are world renowned for building high performance, high quality hardware with custom written software, so the initial announcement of the release of XFS has created quite a buzz in the Linux community.

Flirting with open source

IBM's billion dollar bet

Everyone knows who IBM are, and everyone knows how big and reputable a company they are. Over the last two years, the computer giant has been accelerating their interest in Linux, and have committed a staggering $1billion into both software and hardware production and awareness campaigns.

IBM have put their efforts into a variety of areas including making their software applications such as *ViaVoice* available for Linux and supporting Linux across their range of eBusiness servers. They have also established a number of development centres around the world where they have coders involved in Linux-related work, much of which is contributed back.

Although IBM have not open sourced all of their software, they have contributed many things so far to the Linux community, and they are working a variety of elements at kernel and user level. IBM have also formed strategic relationships with a number of distributors which should raise Linux's profile in the enterprise sector.

IBM are porting many of their popular applications to Linux.

Whether or not IBM choose to rely on Linux as part of their business model is yet to be seen, but given the level of resources they are targeting at it, it appears they are taking the idea of open source seriously.

On being multilingual...

We shouldn't let language loyalty get in the way of our real job: innovation.

Within the development community, there are a broad range of people, practices and preferences. As much as someone may sing the praises of MegaCorp's SnazzSoft++ v1.43242, there will be someone to disagree and plot its demise. This variety of interests is both a good and bad thing. Linux has always given users and developers choice, especially in the field of programming languages.

Some of you may not realise this, but there are bucketloads of different languages, and the open minded Linux community has invented even more. An example of just a fraction of these languages include C, C++, Java, BASIC, Ada, Ruby, Scheme, LISP, Eiffel, Perl, PHP. Although my languages of preference are C++ and PHP, I have seen many flamewars over which is best, but with so many available, why not try out a few.

I have experimented with many of them, and the things I have discovered about these languages is primarily how similar they are. Many of the languages have similar functions and facilities, and anyone who is familiar with Stephen Pinker's *The Language Instinct* will know that this is no coincidence. The other thing I have noticed is that no matter which language you use, you are basically going to end up with the same end product. Our purpose as developers is to write functional software that is useful, and surely it doesn't matter which language you use.

Regardless of language preference, developers have one thing in common – the desire to innovate. We need to remember as coders that these languages, computers and big mugs of coffee are merely tools to help us write good software. The important thing is not to zone in at the tool level, but to focus at the creativity level – a task that is not a problem in the Linux community. Together we can step over the divide and really take the world by storm.

Jono Bacon
The founder of UK Linux, KDE developer and all-round nice guy, Jono Bacon is studying at Wolverhampton University.

The Art of Commenting

Developing an open source project? Andew Arensburger has some time and effort-saving advice.

The author

Andrew Arensburger has been writing open source code since the 1980s, before he knew what it was. In the daytime, he plays a mild-mannered system administrator to support his reading habit. He can be reached at arensb+LinuxFormat@ooblick.com.

Rob Pike, in his essay *Notes on Programming in C* (See boxout) says that he tends to err on the side of commenting less, rather than more. There appears to be a school of thought that has taken this one step further, and believes that comments are, at best, a necessary evil, and that good code should be self-evident enough to obviate the need for comments. Bull.

Yes, you should be writing code that's clean enough that you don't need to explain what it does. But the code only tells you what the code does; it doesn't tell you what the code was intended to do, what it ought to do, what it doesn't do, or why it looks the way it does.

Bad Comments

The canonical example of a useless comment is:

```
//i++; /* Increment i */
```

What makes this a bad comment? Quite simply, the fact that it doesn't add to the reader's understanding of what's going on. Compare this with:

```
for (i = 0; i < num_elements; i++) { frob(elements[i]); if
(elements[i] == 0) i++; /* Ignore the next element */ }
```

In the first case, the comment tells you exactly what the code does, but you knew that already from reading the code. In the second case, however, the comment tells you, in human terms, what the statement is intended to accomplish. This is a minor, but crucial difference: the code tells you what the code does; the comment tells you what the code is supposed to do. If you wanted to change the code so that it used a linked list rather than an array, you would know how to translate that statement:

```
for (elem = elements; elem != NULL; elem = elem->next) {
frob(elem->data); if (elem->data == 0) elem = elem->next; /*
Ignore the next element */ }
```

Comments as Section Headers

Take a look at a good reference book. If you wanted to use this book to answer a question, you might start by looking up a key word or two in the index, or by finding a promising chapter in the table of contents. Then you would leaf through the chapter, reading section titles and table captions, until you found a page that is likely to hold the answer to your question; then you would start reading the actual text. Without the section headers, it would take much longer to find the part you are interested in.

In order for a program to serve as its own reference manual, it should contain chapter and section headings: comments that briefly say what the code that follows does. This allows the reader to skim the comments and skip to the part that he's interested in.

API documentation

Every function should be introduced by a comment. This should say, at a minimum: what the function does, what arguments the function expects, what the function returns, what happens in case of error and which assumptions the function makes (e.g., "Signals are blocked" or "first argument is non-NULL")

As a rule of thumb, include in this comment any information that you would like to see in a reference manual. There are packages, such as Javadoc, Doc++, and KDoc, that read these sorts of comments and automatically generate printable or browsable documentation.

This comment is placed at the top of the function so that the code and its documentation are close together. This increases the odds that, if you change the code, you will remember to change the documentation.

Depending on your taste and local coding guidelines, you may want to adopt a strict style, with headings and subheadings, or you may prefer a more relaxed prose style.

This comment is different in intent from the ones inside the function: this one documents the function's API, and represents an agreement between the function and its caller. If you explicitly spell out the assumptions made by the function, you can avoid performing unnecessary error-checking, and sometimes make it easier to find bugs.

Certain types of comments should go at the beginning of a file: if a module implements a particular protocol, this should be mentioned at the beginning of the module. If it is a documented protocol, such as SMTP, include a reference to the documentation (e.g., RFC numbers). You may then assume that the reader has read this documentation, and you can avoid describing the protocol later on.

If it is a protocol of your own devising with no documentation other than your source, then document it here, or at least summarise it, to help the reader understand what follows.

Error-Checking

Real code does a lot of work unrelated to its primary task, such as error-checking, assertion-checking, context-sensitive help, and so forth. This can add significantly to the length and apparent complexity of your code. For instance, in ColdSync (the author's project which can be found at http://www.ooblick.com/software/coldsync/) a single **printf()** statement grew to over 17 lines of code, once error-checking had been added.

In situations like these, it is especially important to leave section-header comments, lest the reader lose sight of the forest for the trees. You should clearly mark what is important and what is incidental.

Write Comments First

Perhaps the easiest way to make sure that your code has useful section header comments is to write them first. Before you write any actual code, write an outline saying what the code will do:

```
int authenticate() { /* Find out which authentication method to use */ /* If it's a network connection, authenticate host */ /* Prompt user for password */ /* Verify supplied password */ /* If it doesn't match, raise the alarm */ }
```

Then, when you come to actually flesh this function out with code, your outline comments become section header comments.

Writing such an outline carries an additional benefit: it allows you to catch, at an early stage, problems in the design itself. The human brain is a wonderful thing, but it is also a result of three billion years' worth of 'good enough' implementation, and backwards-compatibility back to the earliest chordates. As a result, it is a giant hack with more than a few quirks.

One of these quirks is that the different parts of the brain don't always work together. You may have experienced 'confessional debugging', in which you ask a coworker for help with a problem, but the act of articulating the problem into words, or drawing a graph on a whiteboard, suggests a solution to the problem. The part of your brain that sees the code is only a few neurons away from the part that can fix the problem, but the shortest path between them often leads through the speech centres, out of your mouth, and back in through your ears.

Writing an outline is a variant on confessional debugging: by writing out a compact, high-level summary, you ensure that the design is good, and that you've left nothing out.

XXX

Often, when the creative juices are flowing and you're churning out code as fast as you can type, you'll think of something that needs to be done in the production release of the code, or in the next version, or just a nifty feature that it would be nice to have.

In these cases, stop and add an **XXX** comment:

```
fd = open("myfile", O_RDONLY); /* XXX - Error-checking */
```

For one thing, this tells anyone reading your code that it's still unstable, and also points out where the known problems are. For another, if you don't mark the problems now, they'll be a harder to find a week or year from now, when you revisit your old code.

Open Source Projects

If you are writing an open source project, you expect people to look at your code, submit patches, and help you to improve the project. You should help these contributors. One of the most damning criticisms of the *Mozilla* project was that it was very hard to find one's way around it. Don't make the same mistake.

Why would anyone even consider contributing to your project? In most cases, people just want to make one or two small, specific changes. They want to find the relevant section, fix the problem, and submit a patch. Get in quickly, and get out.

You should encourage these people. How? By making it easy to find the problem. How? Clean code and good commenting.

Section-header comments tell the reader what a given passage does, and allow him to get a feel for the layout of the project before diving into the code itself. Cross-references, e.g.

```
int parse_line(FILE *infile) /* This is used by (*parser->lang)() */
```

help to show how things are connected.

The bottom line is that if your code is too hard to read, people will find it easier to either a) do nothing, b) submit a bug report and expect you to fix it, or c) switch over to some other program that does the same thing as yours, but is easier to hack.

Objections

The main objection to copious commenting, which Mr. Pike raises as well, is that if the comments and code repeat each other, you run the risk that comments and code will drift until they no longer bear any relation to one another. Hence, some conclude, it is better not to comment.

To me, this sounds like saying that a highway construction crew might forget to update the road signs when necessary, so therefore there shouldn't be any highway signs. Of course this is a risk, but in the vast majority of cases, it is nice to have signs that say where the road goes, and roads that go where the signs say.

If you use tables of data in your program, you still have to maintain them, even though the compiler can't tell you whether they're wrong. They're just as much a part of your program as the actual code. So it is with comments. Don't just maintain code. Maintain code and comments.

Commenting Out Code

Just for completeness, I'll also point out that commenting out code is a good way to temporarily delete code that you'll want to return to later. I'll only point out that Perl- or C++-style comments allow you to have comments inside comments:

```
This is a comment # <old code> # old comment
```

whereas C-style comments do not:

```
/* This is a comment <old code> /* Old comment */ this is not commented */
```

Consider this the next time you're designing a language.

Conclusion

Good commenting practices can make code cleaner, easier to understand, easier to debug, and generally more fun to hack. And isn't that why we write code in the first place? LXF

Read on

You can find a copy of Rob Pike's *Notes on Programming* in C at: http://bit.csc.lsu.edu/tutorial/ten-commandments/pikestyle.html

MAILING THE WEBMASTER

Basic CGI scripting in Perl

Perl is great on its own, but marry it to the Web and you have the perfect combination for site automation. Charlie Stross shows you how.

This time round, we're going to look at Perl and the World Wide Web. Perl is the commonest programming language used for automating web sites: we're going to examine a simple example CGI script and see how it uses perl modules from CPAN to get the job done.

Perl came to prominence on the Internet back in 1993–94. At that point, most web documents were static, unchanging HTML files. Where it was necessary to provide access to a database or other program, Perl provided a powerful tool for interfacing to those resources on the UNIX systems that ran most of the web. Today, Apache (with 60% of the web server market) still supports CGI, the common gateway interface, as a mechanism for invoking external programs that produce HTML output.

Our CGI script (see boxout) does two things. If you run it with no parameters passed to it via CGI, it creates and returns an HTML form. It does this the simple way, by printing HTML; there is a more sophisticated way of doing it which we'll described later. If you run the script and populate the parameters specified in the form, it tries to send you email, or rather it sends email to user@host.org – to change this, edit line 12.

We could write this tool as an HTML page and a single-purpose CGI script, such that the script would merely process the results of filling in the form on the HTML page. However, if we do that, we risk having the HTML form direct its output to the wrong place, and we also risk losing track of one half of the application.

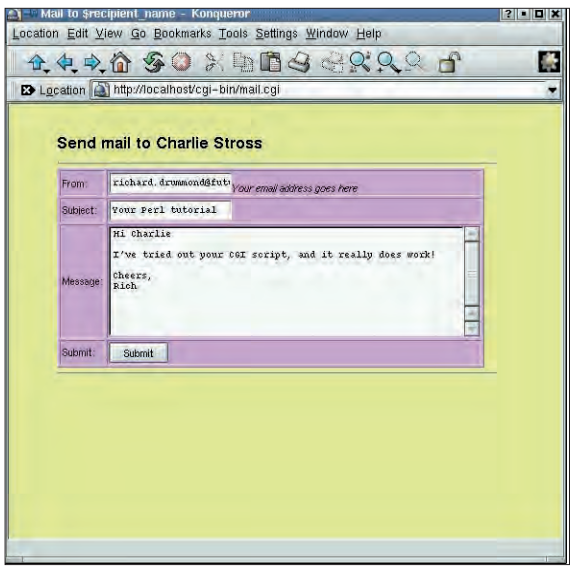

Our example script, sendmail.cgi, creates a simple HTML form for sending email.

As it's a single application which can exist in two states ("fill out form" or "process submitted form"), keeping them bundled together makes installing it easier.

The program starts at lines 07 and 08, which import two Perl modules: CGI and Mail::Send. We qualify the /use CGI/ command with some parameters – these specify the methods we want to explicitly import from the CGI namespace in addition to the normal CGI stuff.

CGI is a big module. If you want to learn about it you really need to read the documentation that comes with it (in POD format) – it runs to about forty pages. Mail::Send, in contrast, is a simple widget designed specifically for a single purpose – to send email to a recipient.

At line 18, we create a CGI object. This implicitly checks our environment to see if we're running as a CGI script, imports bits of data from the web server we're running under, and creates an object we can do useful things with. We do just one thing, at line 20: we look for an HTML form parameter called "live". If it's present, we call a subroutine called process_form(), to process the form input. If it isn't present, we call print_form(), to produce the HTML form for this script. The HTML form contains a (hidden) form parameter called "live". So, if someone's filled out the form (because they've already visited the CGI script), the script knows it's live; if this is a first visit (and they haven't seen the form yet) it knows to print its form out. **print_form()** is the simpler of the two routines, although it doesn't look it at first sight. We pass it a single parameter ($q, the CGI object), then call various CGI methods via $q and print them. CGI.pm supplies tools to let us programmatically generate HTML; call a method like $q->p("some text") and it returns a string consisting of "some text" formatted as an HTML paragraph. **print_form()** could equally well be written as a simple print statement to output HTML. There's no reason for doing this here other than to show off some of CGI.pm's features – but if we were feeling daring we could do nifty Perl things like generate extra fields on the fly, or expand macros. One point to note is that CGI's **startform()** method implicitly creates a form that points back to the current CGI script. So broken links to scripts are a thing of the past.

/process_form()/ does a bit more leg-work. First, we extract the named parameters from the form. Then, in lines 78 to 83 we create and send an email message! This looks deceptively simple; actually, Mail::Internet conceals a vast mass of complexity from us. Just create a new Mail::Internet object (using the **new()** method), use the **subject()** and **to()** methods to say what it's about and who it's addressed to, then call **open()** on it. This

returns a file handle; print the contents of the message into this file handle and close it. When it closes, the mail is automatically sent out.

Finally, we print some HTML – the traditional way, just to demonstrate that you can mix and match HTML and CGI's programmatic HTML generation in the same script, or even the same subroutine.

CGI and Perl

In 1992, when the web was in its infancy, people were just beginning to realise that in addition to serving up HTML documents, web servers could be called on to run programs that produced output in HTML. The Common Gateway Interface was invented by NCSA (from whose web server Apache is descended) as a mechanism to make it easy to interface external programs to the web.

When Apache or some other CGI-compatible server receives an HTTP request for an executable object (one that it executes instead of simply copying to the user), it puts a chunk of information about the request in the environment – an area of shared memory accessible to child programs. The environment is used to store data which can be passed to child processes when the current program executes them; it's also used by the UNIX shells to stash program variables.

When a web server receives an HTTP request for a CGI program, it runs the CGI program, passing it any data from the client by stashing it in the environment. The standard output filehandle of the CGI program is connected to the parent web server process; anything the CGI program prints is read and returned direct to the web client that sent the request.

CGI programs usually need to start their output by stating what type of data they're producing; it's common to see hand-rolled CGI programs start by printing:

```
Content-type: text/html
```

followed by a blank line. This is the minimal HTTP 0.9 header required for an HTML document. If it's missing, the web browser simply won't know how handle the output stream. Indeed, the web server may also complain if it doesn't recognise the output content-type.

There are two distinct HTTP request types that are used to invoke CGI programs: GET and POST. A GET request, like a standard HTTP request for a file, can be accompanied by some »

sendmail.cgi

A fairly simple script for sending email from within an HTML page.

```
01:  #!/usr/bin/perl
02:
03:  #————————————- standard CGI skeleton ————————————
04:
05:  use strict;
06:
07:  use CGI qw(shortcuts font table td TR);
08:  use Mail::Send;
09:
10:  # configuration bits
11:
12:  my $recipient      = "user\@host.org";
13:  my $recipient_name = "Your name goes here";
14:  my $sender_addr    = "nobody";
15:
16:  # end of configuration bits
17:
18:  my ($q) = new CGI;
19:
20:  $q->param('live') ? process_form($q, $recipient, $sender_addr) : print_form($q);
21:
22:  exit 0;
23:
24:  #————————————- support routines ————————————
25:
26:  sub print_form {
27:      # we have no CGI parameters, so print a form
28:      my ($q) = shift @_;
29:      print $q->header();
30:      print $q->start_html(-title => 'Mail to $recipient_name' ,
31:                           -BGCOLOR => '#FFFFA0');
32:      print $q->blockquote(
33:          $q->startform,
34:          $q->p(""),
35:          $q->h1("Send mail to $recipient_name"),
36:          $q->hr(), "\n",
37:          $q->table( {"border" => "1",
38:                      "bgcolor" => "#FFA0FF",
39:                     },
40:                     $q->hidden( -name => "live", -value => "1"),
41:                     TR(
42:                        td("From:"),
43:                        td( $q->textfield("my_name"),
44:                            $q->i("Your email address goes here") )
45:                        ),
46:                     TR(
47:                        td("Subject:"),
48:                        td( $q->textfield("subject") )
49:                        ),
50:                     TR(
51:                        td("Message:"),
52:                        td( $q->textarea(-rows => "10",
53:                                         -cols => "60",
54:                                         -name => "content")
55:                            )
56:                        ),
57:                     TR(
58:                        td("Submit:"),
59:                        td( $q->submit() )
60:                        )
61:                  ),
62:          $q->hr(),
63:          $q->endform,
64:      );
65:      print $q->end_html();
66:      return;
67:  }
68:
69:
70:  sub process_form {
71:      # we detected CGI parameters, so do something with them
72:      my ($q, $recp, $from) = shift @_;
73:      my ($subj) = "CGIMAIL: Subject: " . $q->param("subject") . "\n";
74:      my ($body) = "Originally from: " .
75:                   $q->param("my_name") .
76:                   "\n\n" .
77:                   $q->param("content");
78:      my ($m) = new Mail::Send ;
79:      $m->subject($subj);
80:      $m->to($recp);
81:      my ($fh) = $m->open;
82:      print $fh $body;
83:      $fh->close ;
84:
85:      print $q->header();
86:      print "<HTML><HEAD><TITLE>Mail to $recipient_name</TITLE></HEAD>",
87:            "<BODY BGCOLOR='#FFFFA0'>\n";
88:      print "<BLOCKQUOTE>\n";
89:      print "<H2>Mail sent</H2>\n",
90:            "Your mail has been sent. Thank you.<P>";
91:      print "</BLOCKQUOTE>\n";
92:      print "</HTML>\n";
93:      return;
94:  }
```

variables in the form of name/value pairs (specified in the URL after the address of the script – for example, http://www.foo.com/cgi-bin/myscript?colour=red&option=1). A POST request is more complex: it specifies that the parameters are appended to the request in a MIME attachment. POST parameters can be larger than those of a GET request, can include binary data, and are relatively immune from users tampering with the parameters by manually editing the URL: they're therefore recommended for any non-trivial application of CGI programming.

A minimal – but useful – CGI program looks like this:

```
01: #!/usr/bin/perl
02:
03: use Data::Dumper;
04:
05: print "Content-type: text/html\n\n";
06: print "<HTML><HEAD><TITLE>Current Environment
</TITLE></HEAD>\n";
07: print "<BODY><H2>Current Environment</H2>\n" ;
08: print "\n<PRE>\n";
09: print Dumper \%ENV;
10: print "</PRE>\n";
11: print "</HTML>\n";
```

This script doesn't do anything with its inputs, but it lets you see the environment variables the CGI script can see. To do this, it calls the subroutine Dumper (imported from the module Data::Dumper), to dump out a neatly-formatted version of the hash %ENV. %ENV is the current perl process's view of its environment.

Line 5 of this script prints "Content-type: text/html", followed by a blank line. This is a minimal HTTP header, and is passed through the Apache server to tell the user's browser that what follows is HTML. The following lines print brief HTML headers and then dump the contents of %ENV (at line 9) in a preformatted block.

If you experiment by writing an HTML form, with the SUBMIT button pointing at this script, you'll be able to see the HTTP request showing up in the environment in the form of a variable called QUERY_STRING (if you use the HTTP GET method in your form). For example, if you saved the script (above) as cgi-bin/dumper.cgi (in your web server's file tree), you can write a form like this:

```
<HTML>
 <HEAD>
        <TITLE>Test form</TITLE>
        </HEAD>
 <BODY>
        <H1>Test form</H1>
 <P>
 Enter some data:
 </P>
 <FORM METHOD="GET" ACTION="cgi-bin/dumper.cgi">
<P>
        Favourite colour: <INPUT TYPE="text" NAME
="colour">
 </P>
 <P>
      Preferred food: <INPUT TYPE="text" NAME="food">
 </P>
 <P>
        <INPUT TYPE="submit" NAME="Press me" VALUE=
"Press me">
        </P>
        </FORM>
 </BODY>
</HTML>
```

The result of hitting this form's submit button should be a dump of the environment variables visible to your program, and QUERY_STRING should contain something like:

```
colour=red&food=tomato%30soup
```

(Note that non-alphanumeric characters may be encoded.)

Having to parse this mass of data by hand is a nuisance, which is why we have the excellent CGI.pm module, written by Lincoln Stein (and a host of contributors). CGI.pm takes a lot of the legwork out of writing CGI scripts by letting you treat an HTML form as an object. You can query the object for the current values in a named field, or you can tell it to print out various HTML form elements in a programmatical manner, setting various options along the way. For example:

```
print CGI->textfield("fred")
```

prints out an HTML text entry field within a form, named "fred". And if you're processing data submitted from this form:

```
my $cgi = new CGI;
my $freds_value = $cgi->param("fred");
```

The **new()** method in CGI parses the CGI environment variables and extracts data; thereafter, we can query the value of the CGI parameter "fred" via the **param()** method.

CGI.pm – what it provides

CGI.pm, written by Lincoln Stein and a host of extras, is one of the most powerful Perl modules on CPAN.

CGI.pm defines a class that provides an abstract interface to a CGI session. That is, once you create a CGI object (in a script running on your web server) you can either tell it to emit some HTML in the direction of the user's browser, or process data submitted by a browser. In any event, you don't deal directly with the gory details of the Common Gateway Interface.

A script that uses CGI.pm sits somewhere in an Apache (or other) web server's file space. The web server must be configured to treat it as executable – either by placing it in a cgi-bin directory (where everything is treated as a script) or by using the ScriptAlias directive to tell the server that it's executable. When a user's browser sends a request for an executable object, the web

More information on Perl/CGI Programming

The primary source of information on CGI.pm is, of course, the manual page — type man CGI or perldoc CGI and prepare to settle in for a long read. The man page is long and complex, and is perhaps best read in accompaniment with Lincoln Stein's notes at http://stein.cshl.org/WWW/software/CGI/. If that doesn't satisfy, Lincoln Stein (who is the primary author of CGI.pm) has written a whole book about it: *Official Guide to Programming with CGI.pm* (pub. John Wiley & Sons, Inc., ISBN 0-471-24744-8).

This isn't the only book on CGI programming in Perl. O'Reilly and Associates publish the excellent *CGI Programming with Perl, 2nd Edition* (Scott Guelich, Shishir Gundavaram & Gunther Birznieks, ISBN 1-56592-419-3). This covers sub-topics including querying relational databases, incorporating Javascript to validate forms before they're sent to the CGI script, dealing with browser caches, dynamically generating graphics, and a load of other useful things. If you learn best with a book, this is probably the one to get.

On the topic of Perl and books, I'd like to recommend *Advanced Perl Programming* (by Sriram Srinivasan, ISBN 1-56592-220-4). Yet another of O'Reilly's excellent stable of Perl books, this one covers a number of advanced programming techniques. In particular, if you're a bit confused by Perl's complex data structures, name spaces and modules, this book will make things a lot clearer. If you think of it as a tutorial book that takes up where *Learning Perl* leaves off, you won't go far wrong. In particular, the sections on networking and object orientation are well-nigh indispensable. Published in 1997, this book is likely to require an update when Perl 6 comes out — but in the meantime, it is extremely well-thumbed and occupies a spot on the shelf as close as possible to my desk.

server doesn't simply load the file and send it back to the browser: it runs it, and passes any data the user sent it via CGI.

Your typical CGI.pm script starts by creating a single CGI object. The process of calling **CGI->new()** implicitly parses the CGI parameters and loads them into the object. There may not be any, if for example, this is a straight HTTP request for the script, and it's expected to serve up its own invocation form, but if there is it will be accessible via the **param()** method.

For example:

```
use CGI;
my $session = CGI->new();
my @variables_from_form = $session->param();
```

In this case, @variables_from_form returns the names of all the INPUT fields in the form that was submitted to this script.

If you're looking for a specific variable (say, one called "colour") you can invoke the **param()** method in scalar context:

```
my $requested_colour = $session->param("colour");
```

And you can even modify the stored value of "colour":

```
$session->param(-name => "colour" , -value => "turquoise");
```

Note that in the above example we're using named parameters: CGI is written so that you can effectively pass through hashes of parameters to its methods, making it clear what you want to do. This will be familiar to anyone who's looked at Python's extremely cool parameter passing model (guess where CGI.pm borrowed it?) If you want to do away with the hyphens that indicate a parameter name, call **use_named_parameters()**:

```
$session->use_named_parameters();
$session->param(name => "colour",
        value => "green");
```

or

```
$session->param( name => "list_of_foods",
        value => ["spuds", "peas", "beans", "tuna" ]
        );
```

(which assigns an anonymous array of values to list_of_foods.)

You can also delete parameters:

```
$session->delete("colour");
```

Most web applications track a user through multiple screens (HTML forms) and do different things with data they input in each screen. You can save a user's CGI session data in a file, and restore it later, using CGI.pm's cookie-handling abilities to keep track of where the data they've already entered is saved. This therefore lets you keep track of data entered in earlier forms, providing the illusion of interacting with a graphical application rather than a script that is run like a batch process (i.e. it's given some data, digests it, spits out an answer and exits).

The save(FILEHANDLE) method saves the state of the current CGI object into FILEHANDLE, and you can restore the saved variables into a new CGI object like this (if, say, FILEHANDLE was associated with a file called test.out):

```
        open (IN,"test.out") || die;
        while (!eof(IN)) {
                my $q = new CGI(IN);
                print $q->param('counter'),"\n";
        }
```

In addition to providing access to CGI parameters, the CGI.pm performs another valuable task: it lets you generate HTML. Rather than having to hard-code HTML pages which are sent to the user's browser, you can write code like this:

```
        print start_html(-title => "CGI Results",
                -style => { "src" => "/styles/results.css" }
                ),
        h1("CGI Results"),
        p("Your search returned no results."),
        p("Click",
        a({ href => $q->self_url() }, "here"),
        "to search again"
        ),
        end_html;
```

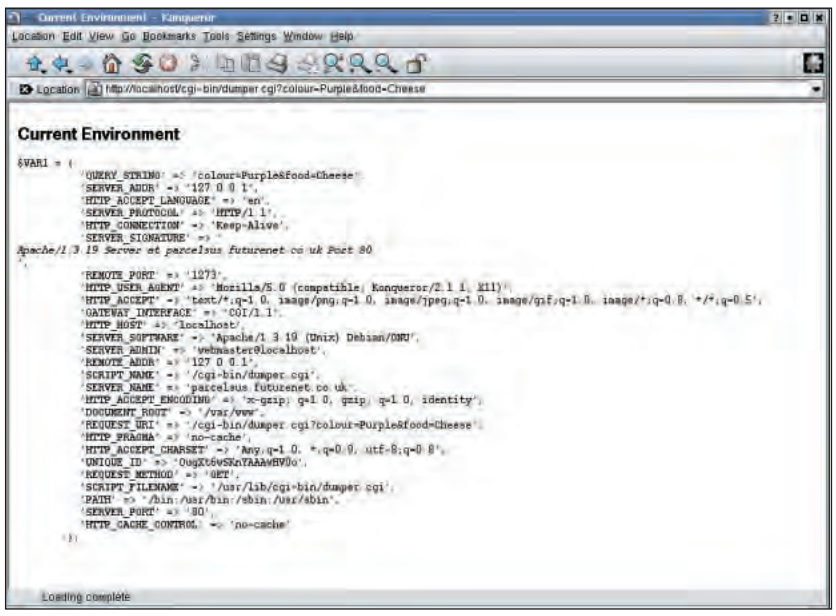

Dumper CGI prints the environment variables that are visible to it.

The **start_html()** method is used to emit an HTML header, and to set up various characteristics of the document (such as javascript to embed, style sheets to use, and so on). **end_html()** is the corresponding end of document method to call. In between, we see **h1()**, **p()**, and **a()** methods, each of which is equivalent to an HTML tag (<H1> headers, <P> paragraph tags, and <A HREF> hyperlinks).

Among the methods we can use are ones for generating form fields:

```
        my $q = new CGI;
        print $q->startform(),
                $q->p("What is your favourite colour?",
                $q->textfield("colour")
                ),
                $q->p("How old are you?",
                $q->radio_group(
                        -name => "age",
                        -values => [ '0-10', '11-20', '21-40', '40-80' ],
                        -default => '21-40'
        )
        ),
        $q->submit(),
        $q->endform();
```

Note how we can mix form elements with the procedural HTML generation described earlier.

The true power of this mechanism becomes clear if you think in terms of using Perl's **eval()** mechanism to take a string of perl and execute it on the fly. Combined with a few simple HTML templates, you can use the ability to generate HTML to make your scripts much more flexible in the way they present search results or forms. For example, you can generate HTML tables to hold your results – and vary the width of the tables to match the largest number of fields returned from a search that returns variable-length records.

» Next month

We take a look at mod_perl, techniques for loading CGI programs directly into the *Apache* web server and making them run faster. We'll also take a look at eval() and error trapping.

FORTÉ FOR JAVA

Speaking Java

Richard Drummond takes you through your first step using Sun's acclaimed Forté development environment.

To take advantage of the fact that Linux Format has brought you Sun's *Forté for Java* – a powerful visual development environment – and the full release of Sun's JDK1.3, Standard Edition on this month's coverdisc, we thought we'd postpone the intended topic of discussion for this issue's tutorial and concentrate on getting you up and running with *Forté*. So, apologies to those that wanted the ins and outs of Java's 2D graphics APIs, but we think you'll find the time spent learning *Forté* will not be wasted.

For full instructions on installing the JDK and *Forté*, please turn to page 92. If you have any problems with installation consult the websites listed in the box on the left for advice.

Running Forté

To start *Forté* you should be logged in under your normal user account. The first time you run it, it creates a workspace in your home directory called 'forte4j_user' to store your projects and settings and also copies some examples projects in here. If you start it as root, as well being a general security risk, any configuration changes you make will be stored in the Forté install directory itself and so will be applied globally.

To start *Forté*, use the command: **runide**.

If you installed Forté in /usr/local – which is the default – then this command should be in /usr/local/bin. On some distros you may need to add this path to your command path. Do this with:

```
export PATH=$PATH:/usr/local/bin
```

In most cases, *Forté* will need to be told where your Java runtime is installed. You do this by setting the environment variable JAVA_PATH. So, assuming you installed Sun's JDK1.3 from our coverdisc, enter:

```
export JAVA_PATH=/usr/local/jdk1.3.0_02
```

Forté should also work with other vendor's implementations of the JDK1.3, such as those from IBM or Blackdown. If you are using a different JDK, then you will need to set JAVA_PATH accordingly. See the documentation for your JDK. *Forté* checks when it starts whether it is compatible with the JDK installed on your system.

If you use *Forté* regularly, it is worthwhile adding the variables you need to your BASH startup scripts. So, for example, add the line:

```
JAVA_PATH=/usr/local/jdk1.3.0_02
```

to the end of the file .bash_profile in your home directory.

Welcome

The first time you fire up *Forté*, you will be presented with the welcome screen.

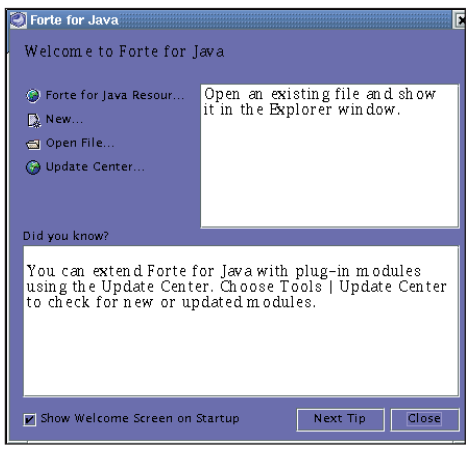

Forté's friendly introduction for new users.

This provides an easy-to-use starting point for users new to the program. If you don't want this displayed every time you start up *Forté*, uncheck the box in the lower left hand corner.

The welcome screen contains four options. The first, 'Forté for Java Resources', opens *Forté's* built-in web browser at the developer resource page on the *Forté for Java* website. Check this site out for new and articles concerning Java programming with Forté. The second option, 'New', lets you create new objects such as Java classes, applets and forms using *Forté's* built-in templates and wizards. More on this late, however. The third, 'Open', lets you open an existing file for editing in Forté.

There's a built in browser should you need to check any developer resources on the web.

Update difficulties

The final option is 'Update Center'. Here, *Forté* can check online for updates to core *Forté* modules, report them to you and automatically download and install them. The Update Center also gives you the option of installing modules that you have downloaded manually. It is possible to get *Forté* to check for updates automatically at given intervals. The first time you start Forté, in fact, an update check is scheduled. If you want to switch this feature off, then select 'Show Autocheck Configuration' and set the period to 'Never'.

There is currently a bug in *Forté's* update feature, however, which Sun are aware of. While you can download updates as normal user, you won't be able to install them, since you won't have permission to write to *Forté's* install directory. To install updates, you will have to run *Forté* as root.

There are obvious security risks to running anything as root, but at present this is the only option if you want to make use of Sun's update service.

On the web

java.sun.com
Sun's Java pages. Contains loads of articles, tutorials, and resources for Java development.
java.sun.com/j2se/1.3
Information and downloads for the Java2 platform.
www.sun.com/forte/ffj/
Sun's Forté for Java homepage.

Find your way around

Forté's main screen looks complicated. Just follow the numbers for enlightenment.

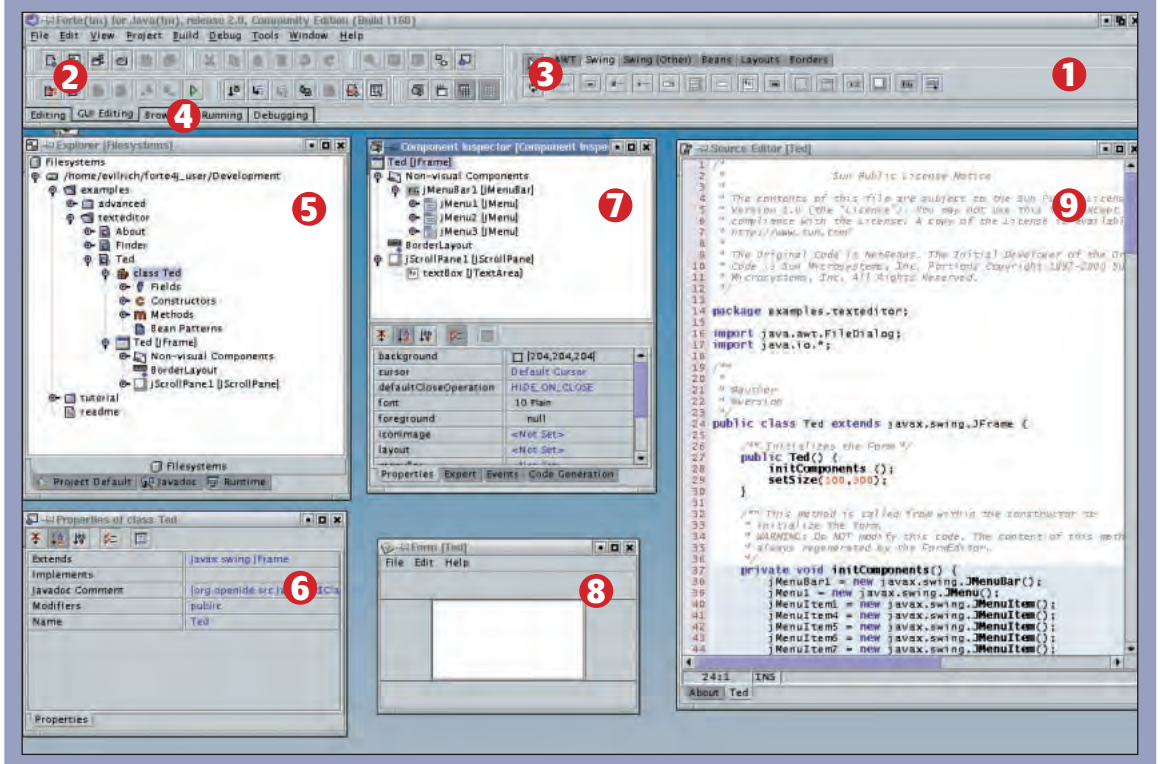

1 Main Window

This window at the top of the screen is *Forté's* control centre. It contains menus, a series of toolbars, the component palette and workspace buttons.

2 Menus and toolbars

Here you can access all the most commonly used of *Forté's* commands. All the controls are context-sensitive, so certain menus options and buttons will only be active when you have an appropriate object selected in the Project Explorer.

3 Component Palette

This contains the graphical components for building GUIs with Form Editor. The tabs split components into groups: AWT, Swing, Swing (Other), Beans, Layouts and Borders. You can add additional components and tabs to this palette if you wish.

4 Workspace buttons

Forté has many windows, so to control the clutter on screen it uses a system called Workspaces. Each workspace corresponds to a particular phase of development and opens an appropriate set of windows on screen. So, for example, the GUI editing workspace puts up the Project Explorer, Properties Editor, Form Editor and so on, but hides the Output and Debugging windows. The Debugging workspace does the opposite. You can choose which windows are shown in which workspace by choosing the Settings options from the Project menu and selecting Workspaces, and you may define your own workspaces.

5 Project Explorer

This shows a tree of files and objects and is split into four panes, selected by the tab gadget at the bottom. The Filesystem panes shows objects stored on disk. By default it shows the work area in your home directory, but you can explore additional directories by right-clicking on filesystems and choosing to Mount Directory. The Explorer lets you transparently browse JAR and ZIP archives, too. The project pane shows which objects from the filesystem tree are part of which current project. The Javadoc pane shows any available Javadoc files, and the Runtime pane shows a tree of running processes, threads, watched variables, and breakpoints.

6 Properties Editor

Here you can modify selected properties of the current object in the Explorer. The options here are context-sensitive to the type of object selected, but, for example, with class objects you can set a class's name, superclass, access modifiers and which, if any, interfaces it implements.

7 Component Inspector

When editing a form, this shows the hierarchy of objects making up your form. The lower half of the window lets you modify the properties of the currently selected object in the tree. The properties are split into four panes: Properties, Events, Layout and Code Generation.

8 Form Editor

The Form Editor lets you design frames, dialogs, applets, etc. You choose components from the palette in the Main Window and place where you want them on the form with the mouse. Thus, you create containers, choose a layout manager, and populate your container with widgets. The properties of the currently selected component will be shown in the Component Inspector.

9 Source Editor

Here's where the real meat of your project is: the Source Editor. This shows the source code containing the currently selected object. With *Forté*, however, you don't actually have to write that much code by hand – the visual tools in *Forté* generates it for you. Since you can define the structure of class with wizards, Forté is able to create skeleton code for you – all you have to do is fill in the method bodies. With the form designer, it also generates code to build and control your GUIs. To help navigate and read the code, stuff that *Forté* has created is highlighted with a blue background.

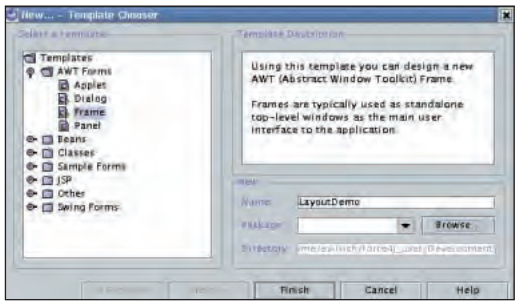

Templates speed up the creation of classes.

First steps

Forté has detailed documentation which you can access via the Help menu and ships with some excellent tutorials to get you started, so rather than duplicate anything covered there, we thought we'd break you in gently and run through a familiar example: we will use *Forté* to re-implement the LayoutDemo example from last month's tutorial. If you followed that article, then you'll see how much quicker and less prone-to-error using a visual IDE like *Forté* really is – at least, once you get used to it.

Your first project

If you haven't already done so, start up *Forté* and close the welcome screen. The first thing to do is to create the LayoutDemo class, which as you'll remember is a sub-class of Frame. To do this we use the class wizard. Select New from the File menu, and this will pop up the Template Chooser – which provides templates that you can use as the bases for your own classes. Some templates will launch wizards to fine tune the creation of a class.

For this example, select Frame from the AWT group. Enter 'LayoutDemo' as the name of the class. The text widget labelled Package lets you choose which package your class will be a member of, but for this example leave it blank – the LayoutDemo class will be part of the default package. The file chooser lets you pick a directory to store the class's corresponding file in. Here we'll just leave it to the default, the directory forte4j_user/Development in your home directory. Click Finish and you're done. The software will then tell you that the object is not part of the current project and asked if you would like it to be added. Say yes, and the class will be added to default project.

Forté will then create the skeleton code for the class LayoutDemo in the file LayoutDemo.java. You should now see the LayoutDemo file listed in the Project Explorer and the code that *Forté* generates listed in the Source Editor.

If you expand the LayoutDemo tree in the Project Explorer, you'll see the objects that *Forté* has created for us in that file. Under the LayoutDemo file is the LayoutDemo file itself – it's the only class contained in the file – and the form definition for the class. The latter is stored in a file called LayoutDemo.form. A level beneath the LayoutDemo class are the components that make up the class – Fields, Constructors and Methods. Expand these out and you'll see that there are no fields currently declared, but that *Forté* has created us a constructor and the methods exitForm(), initComponents() and main().

Have a look at the source code and browse through the generated code. Anything highlighted in blue you cannot edit, it is created automatically from the class's form definition. You'll see that the supplied main() method (the fact that there's a main() means this class can be executed as a Java application) simply instantiates an object of LayoutDemo class and invokes show() to open its window. The class constructor is simple too: it just calls the initComponents() method. This method is where the real action is. As you add GUI elements to with the Form Editor, all the code generated to create and configure these components will be added here. At the moment, it's rather empty. It just adds an event handler to listen for window close events – which just calls exitForm() to quit the program.

There are three main aspects to building GUIs with *Forté*. You design the layout of the GUI with the Form Editor and this will automatically generate most of the code for you. You then declare any additional fields and methods you might with the Project Explorer to do the actual processing or event handling of the class. And, lastly, you write any remaining necessary glue code by filling in the blanks in the source editor. Easy isn't it?

Building the frame

It's time to start building our GUI. The Form Editor should be displaying an empty frame for the LayoutDemo form. Java defaults to using a BorderLayout layout manager for a frame, which is handy, because that's just what we wanted. Select the frame's BorderLayout object in the Component Inspector and set Horizontal Gap and Vertical Gap to 4 in the properties below. This ensures the regions in our layout are nicely spaced apart.

Now choose List from the AWT pane in the palette gadget and click into the west region of your frame to place a list widget there. The Component Inspector will update itself, showing the new list widget as a child of the layout object and its properties below. Switch to the Code Generation tab, and change Variable Name to 'cardList'. While *Forté* automatically generates variable names for each component you create, it makes your code more readable if you give components meaningful names – just as if you were coding it all by hand.

We now need an event handler for this list gadget, so next select the Events tab. Select the actionPerformed event and hit the pop-up gadget. This will open up a list of

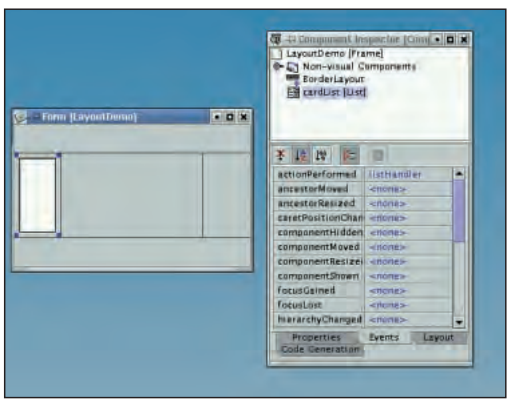

This show the newly added list widget in Form Editor and its properties in Component Inspector.

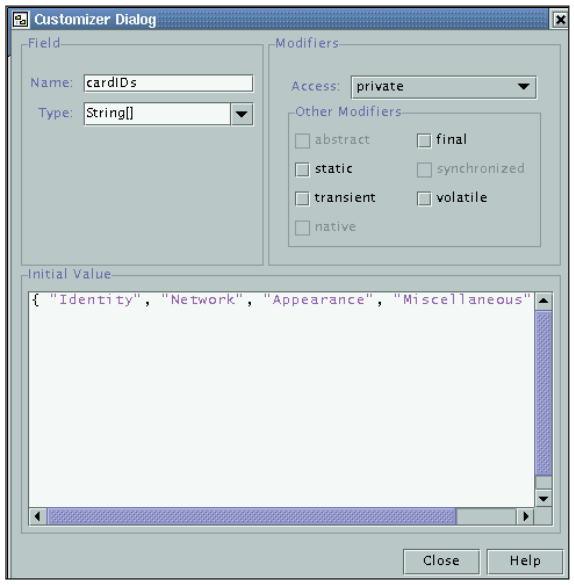

Forté lets you add members and methods to your classes by pointing-and-clicking.

actionPerformed handlers currently defined for our class – which is none so far. Select Add and enter the name of the handler as 'listHandler'. If you remember from last month, we created the handler then as an anonymous inner class, but this will create it as a concrete method. We will fill in the body of the handler in Source Editor later.

Next we need some method of adding entries to the list. To do this we need to create an array of strings. We could do this in the Source Editor but just manually inserting the code we need, but why not take advantage of the tools that *Forté* offers? Go to the Project Explorer and select the Fields entry in the LayoutDemo class. Right-click on it and select Add New Field. In the window that pops up, enter 'cardIDs' for the field's name, 'String[]' for its type, set the access modifiers to private. For its initial value, enter the string:

```
{ “Identity”, “Network”, “Appearance”, “Miscellaneous” }
```

The above procedure will declare a new field as a member of our class – a String array with the above elements – and inserts it automatically in the code. Have a look at the bottom of the file in the Source Editor to check.

Now go back to the Component Inspector, select the list widget and go to Code Generation page. Click on Post Init Code and hit the pop-up gadget. In the window that opens, enter

```
for( int i=0; i<cardIDs.length; i++ )
{
  cardList.add( cardIDs[i] );
}
```

This code adds each item of the array we have just declared as entries in the cardList object. Entering this code as 'Post Init Code' means that it will be inserted into the initComponents() method of the LayoutDemo class immediately after the code that *Forté* generates to create the list widget.

Now for the card panel. Select the Panel tool from the AWT palette, place a panel in the center region of the frame and set its Variable Name to cardPanel. Next we have to create the cards themselves, and, while we could do this manually by using the Form Editor to create and add cards individually, this would be a bit long-winded. Instead we will write the code ourselves – or at least pinch it from last month.

So create another new field in the LayoutDemo class, this time called 'cards' with type 'java.awt.Panel[]', access modifier 'private' and set the initial value to:

```
new java.awt.Panel[ cardsIDs,length ]
```

Now, select the cardPanel object in the Inspector, go to the Code Generation page and enter the following as Post Init Code:

```
for( int i=0; i<cardIDs.length; i++)
{
  cards[i]=new java.awt.Panel();
  cards[i].add( new java.awt.Label( cardIDs[i] + “ Settings”) );
  cardPanel.add( cards[i], cardIDs[i] );
}
```

Did you notice how *Forté* offers code completion options as you type in the names of class library components like java.awt.Panel.

Space is running short, so I'll leave the adding of the 'Help', 'Cancel' and 'OK' in the south region of the frame up to you. The remaining thing to do to make this program a little more interesting is to write the event handler for list gadget. So, find the listHandler() method in the Source Editor and enter the following code in its body:

```
java.awt.CardLayout = (java.awt.CardLayout)
 cardPanel.getLayout();
cardLayout.show( cardPanel, evt.getActionCommand() );
```

That should be it. Let's see if it works.

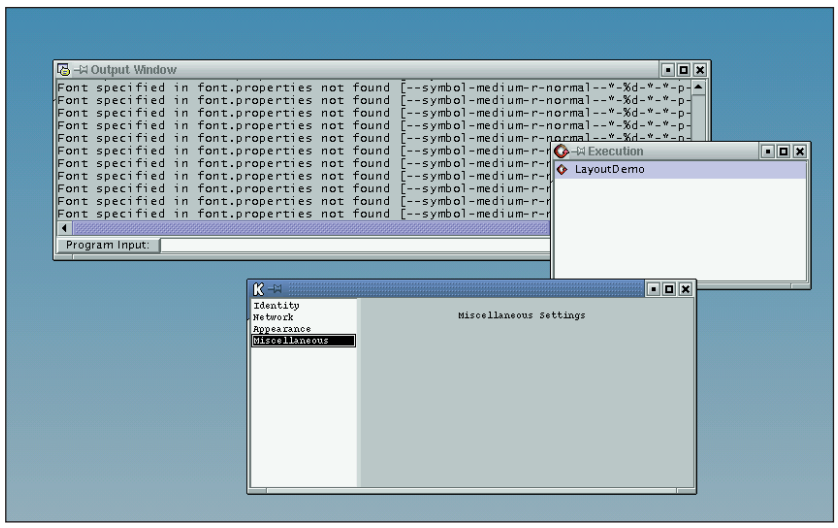

Here's the finished article, and its output, shown in the Running workspace.

Trying it out

You can compile a source file in *Forté* by right-clicking on it in the Project Explorer and selecting 'Compile'; similarly, you can execute it – always supposing it has a main() method – by right-clicking and selecting 'Execute'. Alternatively, if your class is part of the current project and it is set as the main class of the project, then you can select the options 'Compile' then 'Execute' from the Build menu.

When you compile or execute a program *Forté* will switch to the Running workspace, and the Output window will appear. Any compile-time errors will be displayed here and double-clicking an error message will show the corresponding line in the Source Editor. Any console output from an executing application will also be printed here.

Forté has some powerful debugging tools such as the setting of breakpoints, the single-stepping through executing code and the ability to watch the value of selected variables as your code executes. Alas, we don't have to the space to go into all the details here, but Forté's documentation is contains thorough explanations of the program's more powerful features. LXF

LINUXformat Professional

News, reviews and expert advice for Linux pros

THIS MONTH...

Solsoft NP-Lite review **84**

The latest Linux books rounded up **85**

Graham Cluley talks Linux viruses **88**

NEWS

IBM's cleanup campaign

No peace for IBM

While IBM seemed to get a lot of good press for their recent "Peace, Love and Linux" ad campaign, among people certainly not getting the intended message where city officials in Chicago and San Francisco.

As part of the campaign, Peace symbols, a heart and a tux logo were spray-painted onto the sidewalks of the two cities, but sadly not using the non-permanent medium intended. IBM, who have halted the spray-paint campaign, now face a fine and expensive clean-up costs.

Server rivals Sun, who were in the middle of an annual "volunteer week" aimed at helping community projects, couldn't resist offering to help clear up the mess. "For us it's just one more way we're offering to clean up after IBM" volunteered humorous Sun spokesperson Jeremy Barnish.

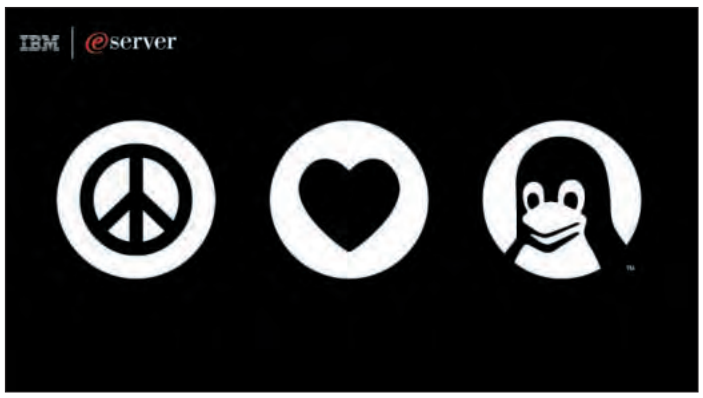

The campaign's good on paper, not quite so good sprayed to a wall.

LXF readers complacent?

Virus could wreak havok

A survey carried out on the *Linux Format* website suggests that an unexpected Linux virus outbreak, such as a more malicious variant of the Ramen Worm, could cause a lot of problems for the Linux community at large.

A huge 34% of people who responded to the survey said they did not see viruses as a significant threat, with a further 25% admitting that they had some concerns, but had taken no precautions. A tiny 12% of over 700 people surveyed claimed to have taken some form of precautions against virus attack.

While there have been few high profile viruses for Linux, and the OS does have better safeguards for restricting potential damage, even the most harmless seeming virus can cause damage. The greatest loss to Windows users plagued by the "love-bug" and "Kournikova" viruses was not in terms of data destroyed, but was more to do with the "bandwidth" theft as mailservers became clogged with copies of the attatchments, the time taken to clear up afterwards and the inevitable loss of face when the virus was passed on to friends, customers and colleagues.

***Linux Format* recently travelled to the infosecurity event at Olympia and spoke with virus expert Graham Cluley. For the full interview and some surprising information about Linux viruses, turn to page 88.**

More deals

Trustix partner IBM on integrated solutions

Norwegian Linux company Trustix are the latest to join IBM as an ISV Business partner. Trustix, who produce their own Linux distribution as well as security and management tools such as *Xsentry* and *Xploy*, will be co-operating on both technology and marketing and will focus on Internet and eBusiness solutions for Linux.

"We are very satisfied with this partnership agreement," said Bjorn Roksvold, Sales Director IBM Norway. "Linux is a major focus area for IBM, and we believe that Trustix has very good solutions for Linux systems management."

www.trustix.no

Comment

Security solution

Rackspace call Exterminator

Managed hosting provider Rackspace have deployed new security features to stave off malicious attacks on their customers websites. The Exterminator Intrusion detection system monitors server centres in Texas and London for any signs of activity which may indicate an attack – if an attempt is suspected, traffic from the possible attacker is halted and information is logged on the attempt.

This is the second new security measure from Rackspace in as many months. In April it released its *Alamo* software to combat a popular tool amongst Linux crackers, Knark "Since day one, the security of our customers' servers has been of paramount importance, and we are always on the lookout for ways to make our systems more secure," said Graham Weston, CEO of Rackspace. **http://www.rackspace.com**

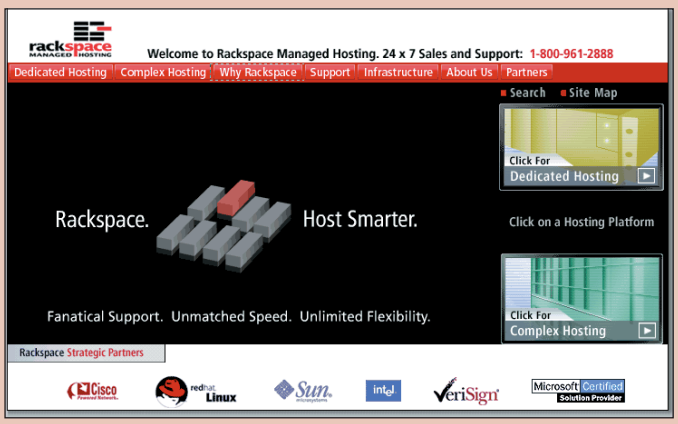

The Exterminator could stop crackers in their tracks.

Shock statistics

One in five may join Linux

The results of a survey carried out on behalf of SuSE suggest that a surprising (even by Linux standards) 23% of computer users will consider switching to Linux the next time they upgrade their computer system.

The poll, carried out by respected market researchers TNS EMNID in Germany, also revealed that over 50% of respondants had heard of Linux, and 10% were already using it in some capacity, either on systems at home or in the workplace.

The key to these numbers seems to be the perceived stability of the Linux OS. 46% of Linux users regard stability as being it's most outstanding asset, as opposed to only 13% of their Windows counterparts. Windows-users were also more likely to associate the term "stability" with Linux.

"The results of the survey indicate that the reasons for not switching to Linux are lack of experience on the part of users and a limited range of applications. Due to wide distribution, Windows still leads in this area" added Wolfgang Best of TNS EMNID.

Testing the testers

Who would you trust with your mission-critical systems?

Choice isn't everything, and because of it, Linux certification is a joke. It means nothing, it adds nothing, and it's going nowhere. Why? Because in the certification arena, everybody is doing their own thing, and few people are in it to make the Linux community more acceptable to business, they are in it to make money for themselves.

Think about it, if someone actually wants to employ, on whatever basis, a certified Linux expert, the likelihood is that the employer themselves know nothing about Linux, or at least not to the extent to test competency. How can they tell the difference between a genuine Linux certification program, and one set up to cash in on the increasing demand for Linux talent – and we have seen more than a few examples of these.

Also, programs tied in to specific distribution vendors fall down, because they fall outside of the open source ethic of choice. It may be good for Red Hat that someone is certified to administer Red Hat systems, but is it necessarily good for the company who is going to rely on that person's expertise?

Now here's an idea. How about companies with a major interest in maturing the Linux industry collaborating on a single, non-profit making organization to administer fair, accurate, non-specific tests worldwide, that everyone could believe in. Tests morever, which people could take without having to spend hundreds of pounds on a whole training course which they may not need.

A good idea, which is why the Linux Professional Institute already exists. But to work, it really needs the support of all the major Linux players. SuSE, Caldera, IBM, Mandrake, LinuxCare, SGI and others are already contributing to this effort – it's time others made the same effort.

Certification is important, but it's more important that whatever certification is developed is as true and fair as possible. The open source nature of Linux presents some special difficulties in this field, but it also presents the opportunity to create a certification system which can uphold the ideals of co-operation and fairness so important to the GNU/Linux community, and become yet another example of why the future is open.

The author of this piece has requested to remain anonymous

SECURITY CONFIGURATION

Solsoft NP-Lite 4.1

■ **LICENSE** Proprietary ■ **PUBLISHER** Solsoft ■ **WEB** http://www.solsoft.com

Security is something we all need to worry about, but can one application simplify your security needs? Richard Drummond finds out.

Solsoft NP is a graphical tool for managing the security policy of your network. It allows you to compile the rules that will be used on any filtering devices on your network (such as firewalls, managed routers, etc.) using a point-and-click interface. *NP-Lite* is a special, free-to-download edition which can manage any Linux firewall and supports the packet filtering features of 2.0, 2.2 and 2.4 series kernels. The full version additionally supports a range of commercial firewalls and routers.

NP consists of two parts – the graphical administration tool itself and the policy compiler – both of which can be run remotely to any filtering device. The admin GUI is implemented in Java, so your administration machine can feasibly be run on any machine on your network with a Java 1.3 run-time. For each filtering device in your network, the compiler takes the policy saved out by the GUI and translates it into a set of rules which can be uploaded to the device. *NP-Lite* can perform the upload of rules itself automatically, which it does – on Linux firewalls, at least – either by a combination of telnet and tftp or with ssh. In the former case, the system logs into the filtering device via telnet and uses tftp to copy the rule scripts – which the policy compiler has created – from the admin machine to the firewall. The GUI itself implements the tftp server during this phase or you can use an existing tftp server.

Mission control

The core of the *Solsoft* admin GUI is the Policy Definition Tool (PDT). This presents a graphical representation of the logical structure of your network. Here, you can add network objects, drag them around with the mouse and modify their properties.

NP defines four types of object: the Policy Enforcement Point (or PEP), that is, any filtering device such as a Linux firewall that is managed by the software; the nexus, any unmanaged filtering device such as a type of router unsupported by NP; the network; and the class.

A network is, as it sounds, a set of machines defined by a set of IP numbers, which can intercommunicate without going through a PEP.

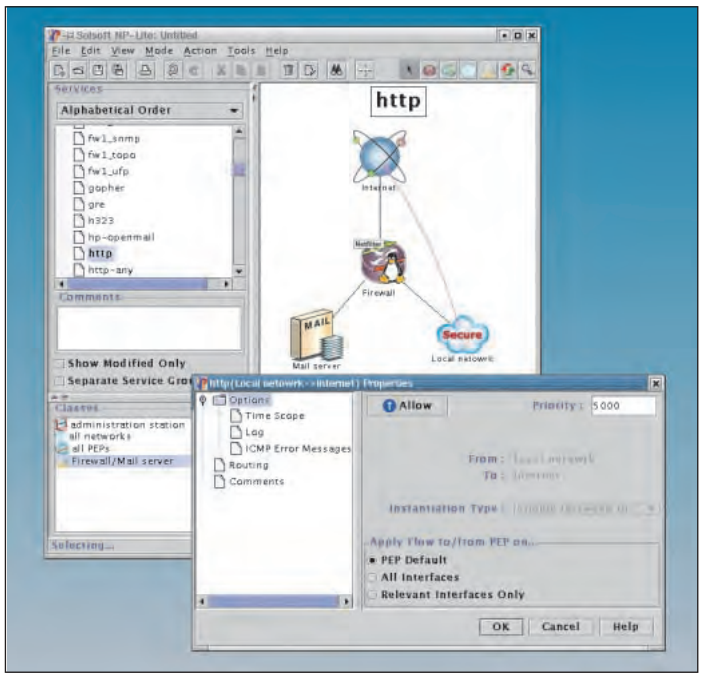

Being able to allow or deny permissions visually makes defining your security policy simple and less prone to error.

> **"Once you have the logical structure of your network mapped out, you can define how and what data will flow around it."**

The class forms a mechanism for grouping together machines with a similar role in the network and assigning them an identifying name. Each machine in a class will have the same security policy applied. So, all the workstations in your accounts department might form one class, since all the users have identical network requirements; IT staff might have different requirements, so the IT workstations would be a separate class. Your web server might be another class, your mail server yet another. It is also possible to define metaclasses and templates.

To modify the properties of an object, you right-click on it and select the properties option from the menu. This will bring up the properties editor, and the options available to you here will depend on the type of the object in question.

In the case of a PEP, you can query its properties via SNMP or configure everything manually. In the Properties Editor you can set its type, specify its network interfaces and the IP numbers assigned to them, the upload method and any passwords required, and set various firewalling options such as error-logging and anti-spoofing globally or per interface. You can also apply a previously defined Network Address Translation (NAT) rule to each interface.

In *Solsoft NP*, NAT rules are created in the NAT editor. NAT is a technique whereby the source, destination or both of a packet are modified as they pass through the filtering device. The editor lets you define NAT rules by name, the type of translation that will be applied and the IP addresses to change.

Permission granted

Once you have the logical structure of your network mapped out, you can define how and what data will flow around it. The first step is to specify which machines are authorised to use which Internet services. The rule in Solsoft is that if a service isn't explicitly allowed then it is denied. Solsoft is shipped with a built-in list of services covering just about everything you'll want to do on your network, but you can always create more if needed in the services editor.

The PDT window shows a list of current services down the left-hand side and the currently selected service is displayed in the top-right hand side. To add a permission to some object for the current service choose the permission tool and drag with the mouse from the source object to the destination object that you want the permission to apply to. Any permission defined for the current service will be drawn as green arrows on the network diagram with the direction of the arrow showing the direction that the flow is allowed. So, for example if you want your accounts staff to be able to browse the web, you drag from the accounts class to the Internet network and an arrow will appear between these two object denoting that flow has been authorised.

The properties of an a permission can be modified in the Permission Editor, which is opened when you right-click on an arrow and select Properties. Each permission has a

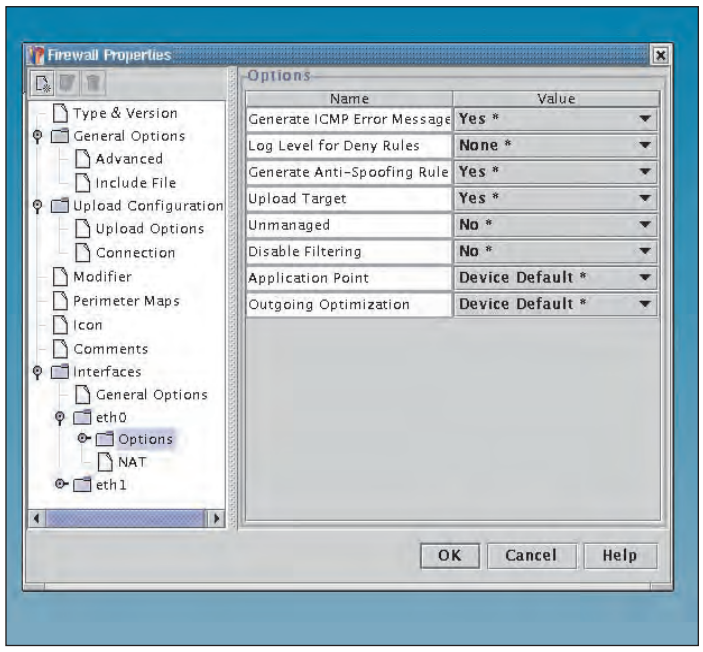

priority assigned to it, with negative numbers given the highest priority. Allow has a default priority of 5000, while deny has -5000. This ensures that the policy compiler denies services rather than allowing them when it comes resolving permissions conflicts when compiling policy definitions for a particular device. You may also fine-tune logging and routing settings for an individual permission and even apply time-restraints. This is great if you want to block certain services at certain times, for example.

Safe as houses

NP-Lite is designed to ease the process of defining and implementing a security policy for a large network. And, to be fair, it does. It has a rather steep learning curve, however. The user interface is not always very intuitive, and although copious amounts of documentation is available as JavaHelp and on the website, I found that it never addressed the questions that I had. Part of the problem is that the Linux-specific aspects of the system are dealt with briefly, if at all, and that the docs require that you learn the language of professional firewalling rather than the cosy Linux terms that you might be used to. Despite these points, *NP-Lite's* graphical configuration offers a much more efficient and less error-prone method of setting up your security policy. The fact that it presents a common interface, no matter what make of filtering device, will be a boon to busy sysadmins.

I do have a couple of complaints. The system doesn't appear to work with networks where the gateway has a dynamic IP address. This is a small point, since NP-Lite is clearly aimed at the enterprise space, but I could envisage smallish businesses with ADSL or cable-modem access to the Internet that could benefit.

There's a load of options to play with in the Properties Editor for setting up your firewall device.

Problems

The other point is that the auto-uploading of rules to a Linux firewall never worked satisfactorily. You can hack it to work, but the rules scripts it generates assume csh rather than bash as the default shell and the upload scripts have a problem gaining root access after the SSH or telnet log in. They need root access to be able to insert kernel modules and call the necessary userspace filtering tools - otherwise failure. You can reconfigure the telnet or SSH server on your Linux firewall to accept root log ins and tell NP to log in initially as root, but this is not something you would really want to resort to on a firewall.

Linux Format VERDICT

Ease-of-use	7/10
Performance	7/10
Features	8/10
Stability	9/10

Makes defining a security policy manageable no matter what size of network you run.

LinuxFormatRating 7/10

BOOKS

SQL in a Nutshell

■ **AUTHOR** Kevin Kline and Daniel Kline ■ **PUBLISHER** O'Reilly
■ **PRICE** £20.95 ■ **ISBN** 1-56592-744-3

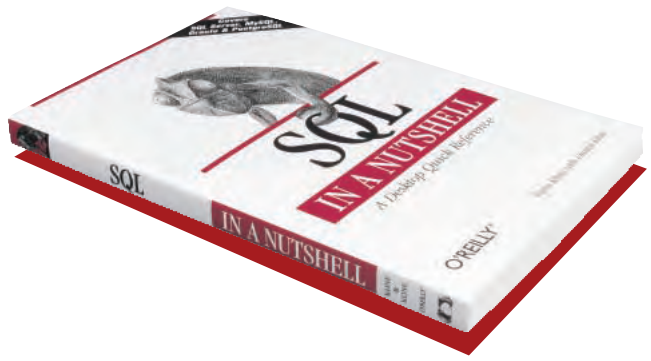

Like the other books in O'Reilly's successful "Nutshell" series this is more of a reference book than a guide or tutorial. It covers SQL (Structured Query Language) as defined by the SQL99 standard and as implemented by four popular RDBMS; two commercial products – namely Microsoft's *SQL Server* and *Oracle* – and the two open source implementations, *MySQL* and *PostgreSQL*.

The first couple of chapters introduce SQL, its history, the SQL99 standard and the concepts of relational databases. The intention of the authors was obviously to offer a concise introduction and set the scene for the rest of the book. But if you don't have, at lease a basic understanding of RDBMS then you shouldn't be reading this book. The space could have been better used, perhaps in providing a few more real-world examples of SQL at work.

The bulk of the book (140 of the 205 pages) is given to Chapter 3 – "SQL Statements Command Reference". This is a very detailed description of each of the statements that make up the SQL99 language. For each statement there is a description of its function and a syntax definition for the standard, and for each of the four SQL dialects. The variations and extensions from the standard by each of the dialects is covered in some detail and illustrated with examples

Chapter 4 covers the various SQL built-in functions. Each is described and examples are used to show the deviations from the standard by the four RDBMs. Four tables cover the functions specific to each of dialects, though these are only lists, so the documentation the particular product would be needed for any detailed information.

The book would be useful for anyone who needs to port a database application from one of the covered dialects of SQL to another, as it is very easy to compare the differences in the syntax between the different dialects.

This is not a book to sit down and read to pick up the basics. If you need to know which SQL command to use, or why a particular command has been used, then you'll be better off looking somewhere else. On the other hand it is a very useful addition to the bookshelf for anyone writing SQL. An experience SQL programmer can check the exact syntax or semantics of a statement, or an occasional user who can refresh their memory of the details of a command.

Steve Heaven

Linux Format VERDICT

A very useful reference book, but not for the beginner.

LinuxFormatRating 8/10

BOOKS

Learning WML and WMLScript

■ **AUTHOR** Martin Frost ■ **PUBLISHER** O'Reilly ■ **PRICE** £24.95 ■ **ISBN** 1-56592-947-0

It seems that O'Reilly books are getting thinner these days. When you compare this with some of their previous tomes, it seems quite svelte. However, it does manage to cover the ground in a reasonable amount of depth, from the basis of WAP communication through to WMLScript (as the title implies).

The book is well laid out with each chapter covering a single topic. The layout works well as it does not try to throw too many concepts at you in a single chapter. Unfortunately, as this quite light on pages, it doesn't quite cover all the topics in the depth that some people may wish for. However, for inexperienced developers, this book could be viewed as the perfect primer to WAP. It will depend upon your level of exposure to Web-based products as to how much use you'll get from it.

The book starts by explaining the concepts of WAP and how communication is achieved from a mobile phone to the Internet and back again. For some reason this explanation is buried in the preface, when it should have a chapter to itself to allow for a more detailed explanation. Then the book starts for real, taking the reader through an introduction to WML and adds to this with further chapter on variables, timers and so on. After 68 pages WMLScript is introduced outlining the various operators, expressions, conditions and so on that can be used to add interest to your WAP site.

By page 158 you are finished (excluding the four appendixes). Amazing, you now know everything about WAP? Well, no, this book does not cover WML or WMLScript in great enough depth, but gives you the knowledge to then build upon. It does work well and would be a good introduction before moving on to a more comprehensive WML book.

My biggest complaint is that there is no coverage of the compatibility problems that exist between different WAP browsers from different manufactures and other such issues at all. This would have been very helpful information for any developer.

Jon Kent

Linux Format **VERDICT**

A good introduction to WML and WMLScript, lacking depth, but will provide a good foundation of knowledge to build upon.

LinuxFormat**Rating**

7/10

BOOKS

Designing Web Audio

■ **AUTHOR** Josh Beggs and Dylan Thede ■ **PUBLISHER** O'Reilly ■ **PRICE** £24.95 ■ **ISBN** 1-56592-353-7

Sound has always been the poor relation in the web designer's arsenal, mainly because of the resources it requires. Serving streaming audio to hundreds of visitors simultaneously puts a great strain on both hardware and available bandwidth and doesn't have the same immediate impact as a well designed flash animation. But as both reduce in cost, and faster connections become available to users, audio is becoming more popular.

Designing Web Audio is aimed at those wanting to add professional quality audio to web sites. It covers the use of sampled sound, as in background loops and sound effects, and streaming audio. It draws heavily on the experience of sound designers from film, television and CD but highlights the differences as well as the similarities, particularly those relating to the limited bandwidth of the Internet.

The text is well written, taking the reader through the various stages of sampling, processing, mastering, encoding and serving audio, while not forgetting that the most important stage in any production takes place before any of these. Deciding what you want to do and who you want to do it for is critical to *how* you should do it. After covering the important aspects of designing and creating your audio, the authors go on to explain the various formats available for providing content over the Internet.

There are separate chapters detailing the use of RealAudio, MP3, Shockwave and Flash, MIDI and Beatnik. RealAudio and MP3 get particularly detailed treatment, including advice on creating audio streams, setting up a server and a discussion of the legal issues surrounding the various uses of MP3. The appendices have a buyers' guide, covering what you need for recording and creating audio as well as serving it over the Internet. There is also a comparison of the various audio formats, detailing the strengths and weaknesses of each, including that most important issue of browser compatibility.

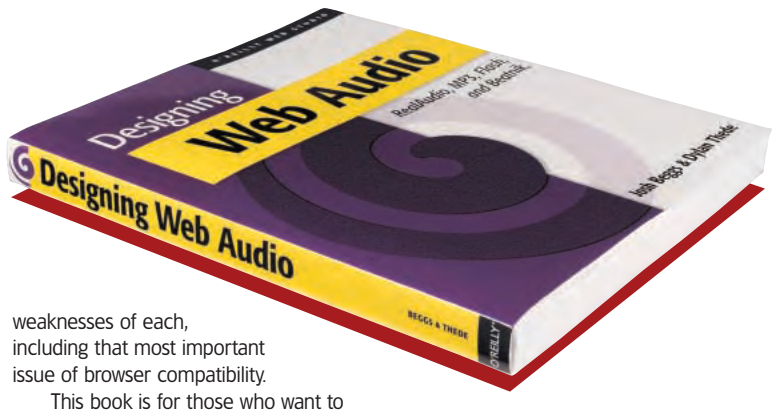

This book is for those who want to add audio to a professional site, either as streamed content or to enhance the overall experience. It is written in a easy to follow style, while providing plenty of information, case studies and food for thought. It doesn't address platform specific issues, although the example screenshots use Windows and MacOS. O'Reilly publish well written, informative books, this one is no exception.

Neil Bothwick

Linux Format **VERDICT**

A good, well written introduction for web developers contemplating adding a little aural sparkle to their Internet projects.

LinuxFormat**Rating**

7/10

No one is immune

Is the Linux virus threat increasing?

Nick Veitch investigates the growing threat of Linux viruses – is it hype, or should we really be scared?

The recent InfoSecurity event, held at London's Olympia, might be the sort of place you'd expect to find *Linux Format* journalists, chatting about Firewalling, Iptables, ssh and so on. But actually we were there to track down something altogether more rare: Linux viruses. And someone who knew something about them. Fortunately our search took us to Graham Cluley, senior technology consultant with anti-virus firm Sophos...

LXF We ran survey about viruses on our website, and the results were probably not surprising. Less than 10% of people were actively taking steps to secure their systems from virus attack. 28% of people don't think there is any need to worry about Linux viruses.
GRAHAM CLULEY To sum up the situation as I see it, there are a handful of Linux viruses, probably the most commonly encountered one is the Ramen worm, and that's probably the only one that has successfully broken out into the wild. As a company, Sophos has been producing anti-virus software for Unix systems for years, even though there was no evidence of a direct threat to these OSes. The reason for that is that Unix systems are commonly used on servers, but desktop machines are likely to be Wintel, so the software checks not only for Unix viruses, but also for Wintel viruses that may try to spread throughout the network as well. That's the major reason for anti-virus software sales on Unix/Linux systems – not only to prevent any Linux virus, but to avoid passing on other viruses.

LXF Are people right to think their Linux box is safe then?
GC Some people say Unix is virus proof. I think that's a dangerous mistake. A lot of people don't seem to understand just what viruses are and how they work. There is a belief that they exist on the Microsoft platform because of security holes in the operating system. In fact of the 64,000 or so Windows viruses in existence, I can only think of half a dozen or so that actually use security exploits to work. All most viruses need to do is be able to copy themselves. Imagine an operating system which wouldn't allow you to copy files – it might be secure, but it would be unusable. Probably the big reason why there haven't been so many viruses for Linux and other Unix platforms is that they haven't grown to the extent where it becomes a worthwhile target. I do believe it's just as easy to write a virus for Linux.

LXF There are more barriers to a virus being able to cause damage on a Linux system though. For example if you are logged on as a normal user, even if you get a virus it isn't going to easily be able to damage or infect files you don't have permission to write too, such as most of the software running on your machine.
GC Okay, but look at the Kournikova virus from recent history. What damage did that actually cause? What it did was send e-mails. Many viruses don't destroy data, they just spread, but that can be just as undesirable. And it would probably not be too difficult to write a virus that understood enough about Linux email clients to do the same. The real damage this caused was in terms of traffic. It caused so much traffic that it brought some servers down, never mind about the credibility factor. You're right, there are more ways under Linux of stopping files without permission, but that's not necessary for a virus to exist. I don't think it's the primary damage that viruses cause any more – it's more about credibility and confidentiality.

LXF There are probably other aspects of Linux that make such viruses less likely to work though. For example the diversity of email clients – there are over a dozen popular ones.
GC Absolutely. The homogeneous nature of the Windows environment is what does it most harm when it comes to viruses. Everyone's using the same word processor, the same email client. With Unix, there's much more diversity. It's just like biological viruses. If everyone's the same and you get a case of measles, you've taken out a population. I don't think it's going to be the end of the world by any means, and you're a lot safer on that platform right now than anywhere else. But there are a contingent of Linux users who think that they are automatically immune and I think that's very dangerous – especially as just this year we've seen more viruses than ever on Unix platforms. Linux is more high profile and more available now, and the virus coders are paying more attention to it.

LXF Do you think that it's simply a matter of the platform's popularity then?
GC It's partly that. It's partly because the virus writers are using Linux themselves. And it's cool. Plus, we see about thirty new PC viruses every day. If you are a virus writer and you want your name or "handle" to get in the papers – well, if you write a Unix virus it's more likely to be on ABC news or ZDNet at least, because it's a novelty. That's why someone wrote a virus for the Palm, even though it's never affected anyone – he knew he'd get headlines and cause a stir.

LXF We spoke to a reader in the wake of the Melissa virus. They said that it was pretty easy for them to filter stuff out at the server level.
GC Some people do that, and it can be effective in the short term, but you have to remember to turn it off again. There can be problems. I know some people wrote scripts to delete any message which said "I love you", but not only did they block the virus, they blocked the updates from the anti-virus companies which said "I love you virus update". There are some simple things I would always recommend. For example, why would you want anyone outside of your organisation emailing you an executable file, or a vbscript? It might be okay from inside the organisation, but coming from outside? These simple precautions can help security features, even if you do have anti-virus software installed.

LXF It interesting, but we see that most Linux users, and most Linux distributions developers, are very aware of the security implications

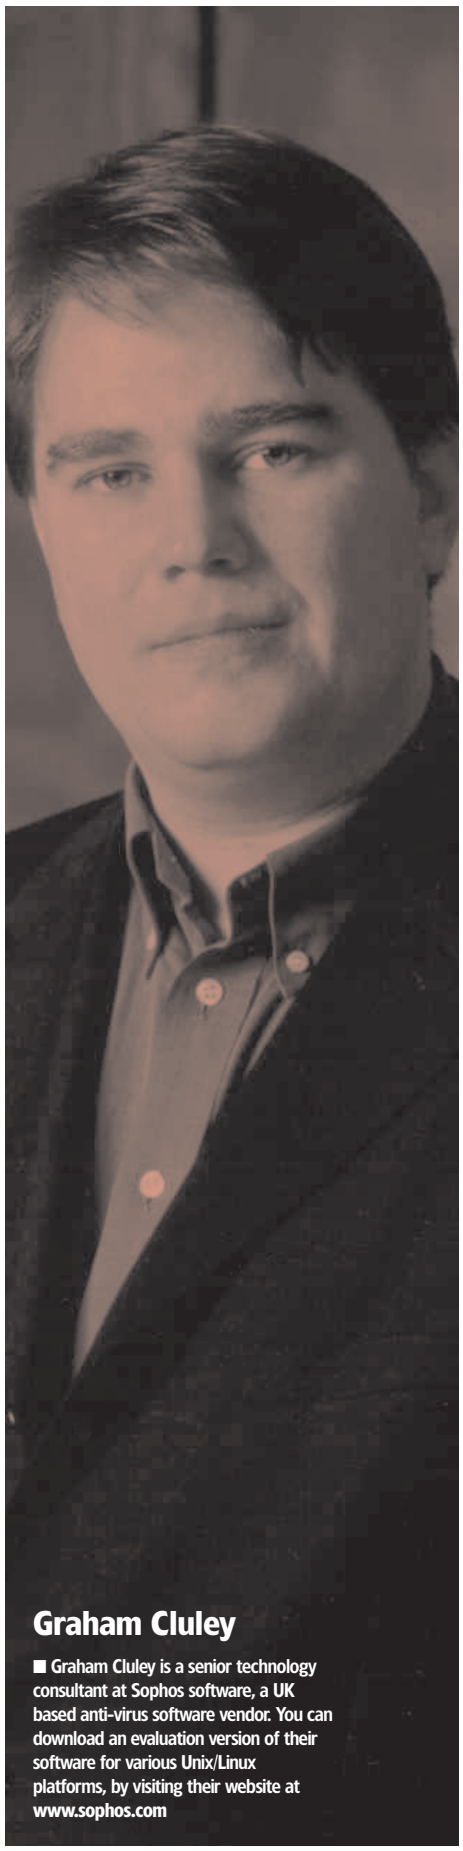

surrounding servers and websites, but there seems to be less concern over developing anti-virus measures.

GC Well, the other thing is that we're seeing a big increase in back-door trojans – the kind of thing that attacked Microsoft recently when someone sent them a trojan horse, which opened a back door to allow them in to look at *Minesweeper* code or whatever – if only they'd destroyed that paperclip thing while they were there – anyway, I think we're seeing more of that sort of activity on the Linux platform. The good news is that the Linux community as a whole are better at applying patches and updates than the Windows community. Of course, as you go more mainstream, you may find that there are more Aunt Agathas or whatever, who are going to be less conscious of security updates, opening up more opportunities in the future.

LXF Are sales of Linux anti-virus software booming?

GC Unix sales are a small part of Sophos's business, and we're aimed more at professionals than home users. However, we do allow companies who have a site license to let their employees use the software for free at home.

LXF That seems like a smart move because a lot of viruses probably get into the workplace from disks people have brought in from home.

GC Absolutely, and potentially more damaging because they can bypass whatever server protection you have stopping stuff coming in from other places.

LXF Are inherently mores secure operating systems like Linux going to have an effect on curbing the devastation wreaked by viruses?

"It's just as easy to write a virus for Linux."

GC Certainly in terms of data destruction. I think increasingly it's the bandwidth, the credibility and the confidentiality issues which will cause the problem. I don't honestly think it is possible to create an operating system that is immune to viruses – at least not a useable one. I'm afraid that viruses are here to stay, for as long as people care to write them. What we'd like is for people like Tony Blair and the government who are pushing the use of computers, and giving more and more kids the opportunity to use them from an early age, to think about teaching computer ethics as well. I think generally anti-virus software does a good job, but we need to educate the users too to be sensible.

LXF Do you think it really is possible to teach the ethics of computers?

GC With kids, I like to be optimistic about it. Just like trying to teach them not to give each other Chinese burns, we can make headway in teaching them not to do it electronically either. It's not just viruses, it's things like privacy too – you would probably think it wrong to read someone's diary, but its just as bad to use someone's computer without their permission.

It's a real shame though. I know kids might think it's cool to be on the "dark side" and get involved with cracking and viruses, but they really have made computing a lot less fun. I started out on a ZX81, and in those days you didn't have to worry. You could load up stuff from tape without concern. Now people are deterred from getting the benefits of computing. Your Auntie Agatha again may be put off from emailing her relatives in Australia because the last time she connected to the Internet she lost all her data or her computer no longer works properly. I think if kids could identify with the victims, then maybe they'd think twice.

LXF Perhaps because Linux comes from a more developer culture – it was conceived and created by developers...

GC You think there are more good guys?

LXF No, I wouldn't claim that, but I think there are certainly more positive role models. People are probably more likely to aspire to be like Linus Torvalds, Alan Cox, Miguel de Icaza or whoever...

GC Right, but it's all more mainstream now – it's the guy in the street with the baseball cap on sideways that wants to run Linux rather than Windows now. And inevitably that means you're going to get more scumbags. In the early days it may have been more altruistic and people were genuinely working towards good goals, and Linux is a very good example of that, inevitably there are going to be more dark horses involved now who are less altruistic.

LXF Do you think the current government, really understands the technology enough to create policies that would help?

GC I'm not confident the ministers do. But the Police do – there is a much better level of expertise there, and they have had success with paedophiles and cracking groups. The trouble there is one of economics. They've had £25M to invest in a computer crime unit, but compared to the money involved in the Internet economy that's peanuts.

LXF Is the future bleak, then?

GC No, there's nothing to panic about. Not a big worry for Linux users right now, but don't keep your head in the sand. LXF

Graham Cluley

■ Graham Cluley is a senior technology consultant at Sophos software, a UK based anti-virus software vendor. You can download an evaluation version of their software for various Unix/Linux platforms, by visiting their website at www.sophos.com

Coverdisc

Feeling chuffed at finally getting Forté for Java and the JDK, Paul Ravening is your guide through the wonders of this month's jam-packed Linux Format CD.

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.

On the CD

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

Read me first

When we get software suggestions put forward to us, either by mail, phone or the forums, we always try to act upon them. The very nature of open source software makes it easy for the most part to supply our readers with the software they want.

However, sometimes a product comes along which is asked for constantly, demands a lot more work to organise, but is definitely worth it in the long run. "What is this?" I hear you cry. Well after what seems an age, we are pleased to bring you Sun's *Forté for Java Community Edition* plus, of course, the Java Development Kit to accompany it.

Sun's *Forté for Java* software is the only development tool needed for building Java applications on any platform. It is fully modular and easily extensible because it's completely open and based on the NetBeans Tools Platform. The *Forté for Java* IDE enables you to easily and quickly create Internet services and solutions with pure Java code. The Community Edition product includes a complete and highly-integrated set of tools – including a Web browser and server. With this, any developer can build stand-alone applications, applets, JavaBeans, and Java clients. Be sure to read Sun's license agreement on the CD though.

GNOME1.4 is finally out, and we have both the source from the GNOME project itself and the Mandrake RPM's for the brilliant version from those clever folks at Ximian. The latest version comes tied together with the excellent *Nautilus* which makes it a pleasure to use and in my view, it's now got the friendliest interface I've ever seen on an operating system. I'll certainly be sticking with it. But enough of me, if you fancy giving it a go, then skip over to the installation pages and get cracking. And if you find yourself at a bit of a tangle

trying to install it, then there's a complete site, grabbed and stuffed onto the CD to help you along your way.

CompuPic is a neat little digital content manager which enables you to view all kinds of picture files, with neat features such as previewing images before opening them, performing alterations to photos without a bulky paint package.

We've also got *Audacity*, an open source audio editor which, although competent, is in it's infancy at present, but is showing great promise for the future. We'll have to wait and see. And we've got the latest from *Mozilla*, release 0.8.2, which brings a whole host of improvements over the last release.

And last but not least, you'll find a couple of Hot Picks, the latest kernel and more...

Paul Ravening
NEW MEDIA EDITOR

LXF15 June Disk contents at a glance...

SUN'S FORTÉ FOR JAVA
This is a real 'you asked for it and we're happy to give it to you' moment! After wading through reams of agreements we can finally bring you – for the first time on a UK Linux magazine – Sun's Java Integrated Development Environment (IDE). Of course it would be no good without the full JDK, so we've thrown than on for you too.

GNOME 1.4
Getting the Ximian version of the latest *GNOME* desktop was almost a "stop the presses" affair. But get it we did! Then we added the GNOME + Fifth Toe suite of applications as well.

AUDACITY
This is a very respectable open source audio editor which got us excited thanks to the prospects of compatibility with the hundreds of VST plug-ins on the 'net. Unfortunately that's not implemented yet. It's still a very competent app though.

ANJUNTA
A rather versatile C/C++ IDE.

PLUS: KDE 2.1.1, Mozilla 0.8.2, latest kernel source, CompuPic, HotPicks, Roundup PIMS and more...

Anjuta

If Java isn't your favourite thing, why not try this IDE for GNOME.

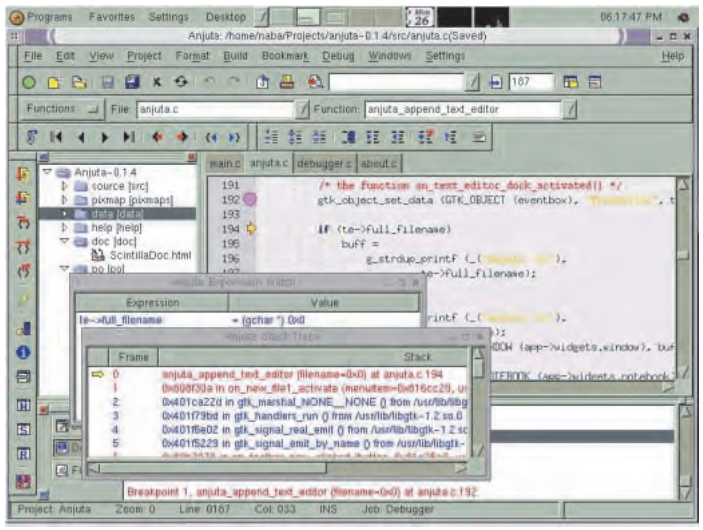

Catching those pesky bugs in *Anjuta*. Note that this IDE has an onboard debugging platform.

***Anjuta* is a very versatile Integrated development environment** for C and C++ (Gnu/Linux). Written in and for gtk/GNOME, it features quite a selection of advanced programming tools and utilities. *Anjunta's* key features include project management, application wizards, onboard interactive debugger, and a powerful source editor with source browsing.
Given the complexity of many of the command based development tools we've come across, the programmers of Anjuta seem to have done a good job of reaching the aim of putting a friendly face onto their app.

Anjuta is an effort to marry the flexibility and power of text based command line tools with the ease-of-use that graphical interfaces such as GNOME can provide. The developers have tried their best to make it as much a 'user friendly' interface as possible, but it's still not something the casual user not versed in the ways of C and C++ (like myself) can just pick up and use. It still looks a little forbidding, but we're sure seasoned code poets will find it a powerful, yet uncomplicated tool.

If you like *Anjuta* and fancy keeping an eye on its development in the future, you'll find all the news regarding the project, including announcements of new releases and updates, at the the project's sourceforge homepage http://www.sourceforge.net/projects/anjuta/.

Audacity

An open source audio editor.

Audio manipulation under Linux has never really taken off, due to the multitude of programs that have gained names for themselves on other platforms. *Audacity* could change that.

This is a program that lets you manipulate digital audio waveforms. In addition to providing facilities for recording sounds directly from within the program, it can import a range of sound file formats, including WAV, AIFF, AU, MP3, and Ogg Vorbis.

Audacity is an excellent editor for preparing samples or loops for use in any of the popular tracker programs. It supports all the common editing operations such as Cut, Copy, and Paste, plus it will mix tracks and let you apply plug-in effects to any part of a sound. Admittedly, the plug-in part of the program is limited at the moment, but the project is coming along in leaps and bounds. It also has a built-in amplitude envelope editor, a customisable spectrogram mode and a frequency analysis window for audio analysis applications.

All you could want in an audio program. *Audacity* is currently in version 0.95. It is not finished yet, and there are quite a few things that the developers want to get right before they release version 1.0.

However, many people have found that it's quite stable enough to use, and its many unique features and intuitive interface make it more enjoyable to use than many other audio editors.

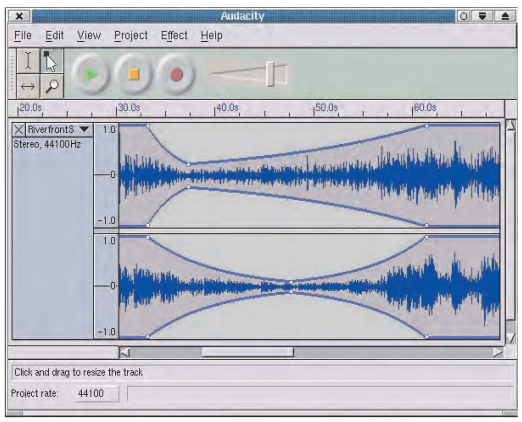

Audacity provides a decent selection of editing tools to help you prepare the perfect sound.

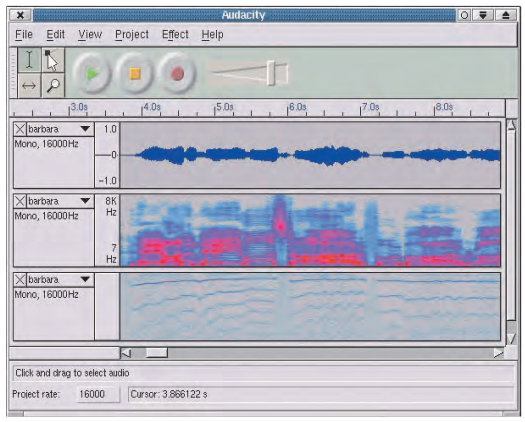

What your sound looks like: these are Waveform, Spectrum and Pitch displays of the same track.

It's your CD

We keep saying it, but, well, it is. And you can help decide what's on it.

You tramp down to the newsagents every month and buy the magazine, so we'd really like to know what you want to see on your Linux Format cover CD (or even DVDs) in the future.

Would you like more complete distros, documentation, games or business apps? Or perhaps you have your own open source project you want to get into the hands of Linux Format readers. Let us know.

As you can see, this month we've worked our socks off, and managed to secure one of the most asked for programs in the history of the magazine – Sun's Forté for Java – and we'd like to continue to give you just what you want.

So don't be shy, reach for that email client and send an email to me at paul.ravening@futurenet.co.uk. You could also post a message on the forums located at our site http://www.linuxformat.co.uk. You could even put pen to paper and send us a letter.

We look forward to hearing your suggestions.

Sun Forté for Java 2 Community Edition

At long last we're delighted to bring you this excellent and most requested Java development tool.

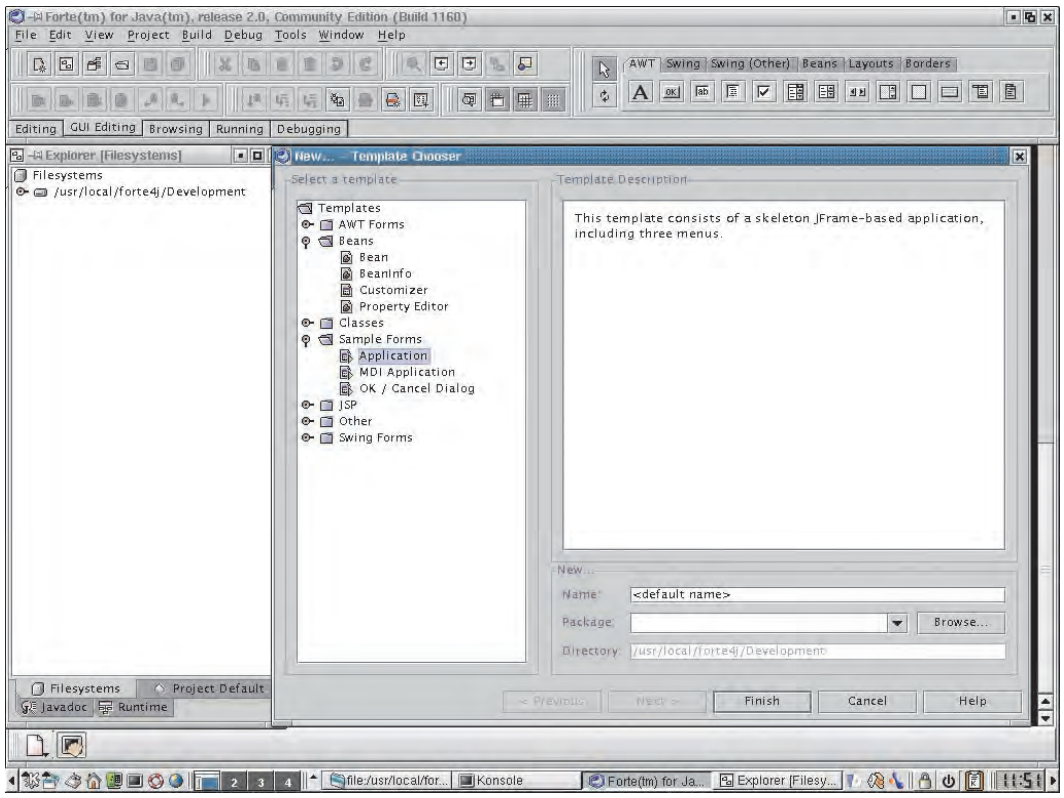

The template editor is the place to create new classes.

Here at *Linux Format*, we're very pleased to be able to include this piece of software at last. And we thought, as it's taken so long to get our sticky fingers on it, that we'd give readers who have no idea what Java is the chance to use it, by including a short primer for the language here. If you're already a Java Guru you might want to jump to the installation instructions on the next page, but if you don't know your class (library) from your elbow, this is the ideal place to start.

So what's all this Java then?

Simply put, Java is a general-purpose, object-orientated programming language originally designed by Sun for embedded applications. Syntactically it looks a lot like C++, but that's where the similarities end. In spirit, it is more like the pioneering object based language Smalltalk.

Java was created with cross-platform portability in mind, so you could use Java programs on a wide range of different platforms. But many different platforms means many different types of language I hear you cry. Well, not with Java. Java programs use an architecturally neutral format. That is, rather then being stored as instructions for a particular processor, Java programs are compiled into and stored as Java bytecode. This bytecode is then translated by an interpreter at runtime. This interpreter is called the Java Virtual Machine or JVM.

The JVM will allow Java programs in their native java bytecode state to run on any platform – and you'll find a JVM for most popular platforms you can think of. As long as someone's created a JVM for your machine, you'll be able to run, and create Java programs. Once this JVM is active, Java from any source – usually websites – can be run on this machine. As well as providing portability by running Java programs in a virtual environment, the JVM offers other advantages such as safety and security.

JVM's are now normally packaged along with browsers to allow you to view Java programs on the Internet. These are normally known as Java applets and are typically, small programs that must be run within the context of the web browser or applet viewer. Applets have been created to do all sorts of tricks than can liven up a static web page, from playing animations and sound to creating interactive games, so is popular with web designers. It has no access to the environment outside of your web browser so it is secure. A rogue applet can't do nasty things to your system.

Three components

There are three key components that you need to program in Java: A compiler, the Java Virtual Machine and the class library. The whole package is called the Java Development Kit or JDK. Only the virtual machine and class library are needed for actually running a Java program and so these two together are called the Java

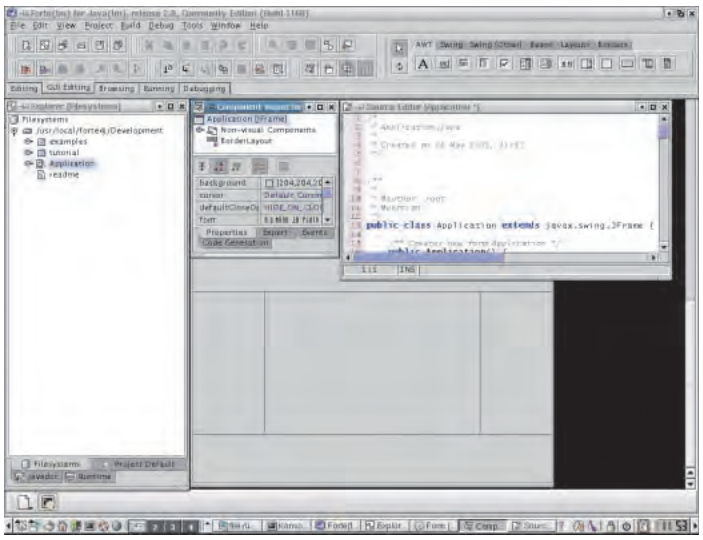

Forté makes light work of creating forms.

Runtime Environment. You'll find a copy of the official Sun JDK on the CD in case you don't have it installed. This is something you'll definitely need to run Forté for Java.

RAD Dude!

To help you along with creating amazing Java programs, Sun have created the definitive development suite in the shape of Forté for Java. Integrated, extensible and easy to use, you'll get started fast with a seamless interface that provides powerful developer tools, such as an application browser, Java language source editor, debugger, compiler, online help and much more. There are also wizards to help you along the way, such as the deployment wizard which enables you to quickly and easily create deployable .JAR files of your own applets, applications and JavaBeans that you've created. The editor itself provides you with a friendly graphical user environment with a form editor to create dialog boxes, windows, applets and more, allowing for true RAD (Rapid Application Development) of visual applications. Your productivity will increase no end!

So have a look at the next page, install Forté for Java and enjoy Java programming development at it's best.

Installing

Forté for Java is a simple program to install if you follow our instructions. The only requirement is that you need JDK 1.3 installed on your machine to make it work – you must have a full JDK not just the Java run-time. We have provided Sun's JDK1.3, Standard Edition on the coverdisc, but Forté will work with JDK1.3 implementations from other vendors such as those from IBM and Blackdown. You can have more than one JDK installed at one time, but you should ensure that only one is present in your command path at once or you will get confused.

Both the JDK and Forté need to be installed with root privileges, so use whichever is your preferred method of obtaining these. Logging in as root is the easiest – but least safe option.

Installing Java 2SDK

1 The JDK is present on the coverdisc as a self-extracting tarball. You just need to execute that file to install it. So CD to the directory that you wish to install it to: typically this will be /usr/local. Then, assuming that your CD-ROM is mounted at /mtnt/cdrom, start the installation by typing:

```
sh /mnt/cdrom/Sun/j2sdk-1_3_0_02-linux.bin
```

2 You will then see a license agreement. It's important that you read this and understand the terms and conditions of use of the software. When you have read this (and presumably agree to the t&cs), type yes and hit return to proceed.

3 The JDK files will now be extracted to a folder called jdk1.3.0_02 in the current directory. If you are installing to ./usr/local/ as we suggested then this directory is /usr/local/jdk1.3.0_02. Remember this directory, since its the home directory of this JDK install. Some software, including Forté for Java, needs to be told the home directory of the JDK it should use.

4 The JDK is now installed and will work fine. If you want to run any of the Java executables from a terminal, you should make sure that your command path now includes this Java directory. You can make sure of this by editing the .bashrc file in your home directory. You'll find a line in this file that says "PATH=/sbin..." Just add :/usr/local/jdk1.3.0_02/bin' to the end.

There is a minimal amount documentation covering the JDK tools can be found in the JDK home directory. For more information see java.sun.com/j2se/1.3.

The IDE comes with an integrated web browser, which is useful for testing your projects.

System requirements

Sun's JDK1.3, Standard Edition, has the following system requirements:

- **Pentium class processor**
- **32MB RAM (48MB recommended)**
- **Kernel 2.2.12 or later**
- **glibc2.1.2 (or later)**
- **A 16-bit X display is also recommended.**

Sun's Forté for Java, Community Edition, has the following requirements

- **Pentium II 500 MHz (800 Mhz recommend)**
- **128 MB RAM (512 MB recommended)**
- **128 MB hard disk space (384 MB recommended)**
- **JDK1.3**

Installing Forté

1 Installing Forté is much easier, since it comes packaged as an RPM file. For distros which use the RPM format, all you have to do is type:

```
rpm -i /mnt/cdrom/sun/ForteCE-FCS-2.noarch.rpm
```

If you are using a Debian-based distribution, then you'll need to convert the RPM file to a DEB before you can install it. CD to a temporary location in which to build the DEB – you need root access, so your /root directory will do just fine. Enter the coim command:

```
alien /mnt/cdrom/sun/ForteCE-FCS-2.noarch.rpm
```

Install the resulting package with the command:

```
dpkg --install fortce-fcs_2_all.deb
```

2 The Forté for Java package will now be installed – it may take a minute or so for the process to complete.

3 Now that the Forté package has been installed, you can run it with the command:

```
runide -jdkhome /usr/local/jdk-1.3.0_02
```

You should always make sure that you run Forté under a user account and not as root.

Alternatively you can edit your .bashrc file and add the lines:

```
JAVA_PATH=/usr/local/jdk-1.3.0_2
export JAVA_PATH
```

Then you can start Forté simply by typing **runide**.

For an introductory tutorial on how to use Forté, please turn to page 78. More information and documentation is available from the Forté for Java homepage at www.sun.com/forte/ffj.

Ximian GNOME 1.4

Paul Ravening clicks the heels of his fetching ruby stilletto heals three times and says "There's no place like GNOME."

Over the last few months, we've had a plethora of important milestone updates. The 2.4 kernel, *KDE 2.1* and now *GNOME1.4*. Yes, finally the *GNOME* developers have churned out the latest version of their acclaimed desktop environment. It's a major step forward for the project, with the new version containing a wealth of new features. Amongst other improvements, *GNOME1.4* is the first version of the desktop to include the much touted integration of the *Nautilus* file manager and web browser from developers Eazel.

We covered Nautilus in detail last issue, concluding that it was both beautiful and innovative, if a touch on the temperamental side.

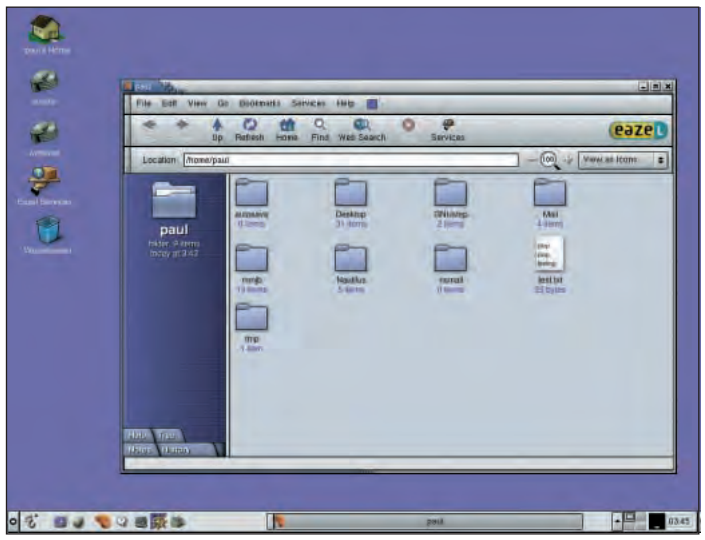

***Nautilus* is integrated into *GNOME* for the first time.**

Installing

Unfortunately, as we were hitting deadline, there were no RPM's in existence – that's how up-to-date we are! If you're unsure about compiling *GNOME* and removing your old version manually – and aren't afraid of a bit of download time – then pop over to http://www.helixcode.com. From there, you can install *GNOME1.4* to your system from the RPM's contained on their server.

But if you're up to the challenge, it's pretty straight forward to compile your own software. Use the standard **./configure; make; make install** command sequence to compile and install each package.

If you choose to specify a —prefix option to the **./configure** command, please specify the same one for all the packages – current *GNOME* limitations prevent full functionality unless this is done. If you are installing from source onto a system that has *GNOME* already installed via the packaging system then, to be safe, you should remove the packages before compiling.

When reinstalling or upgrading a source installation, also back up and remove your home directory and any *GNOME* related directories (.gnome, gnome-desktop, etc.) to prevent possible start-up conflicts that could occur with the new version.
Some systems may also need files from the ftp.gnome.org/GNOME/extra-src/sources site to successfully compile the desktop environment.

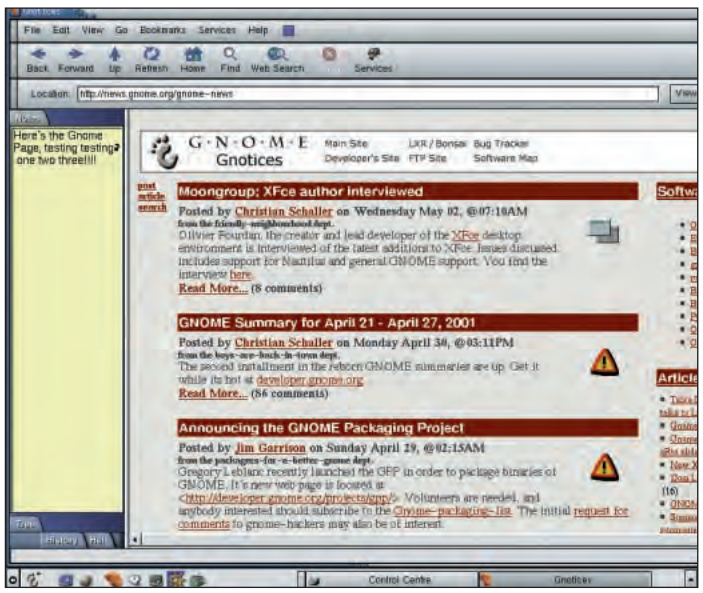

***Nautilus's* notes facility is excellent for jotting things down mid-browse.**

Compiled?

If you are using Red Hat 6.x, and *GNOME* is already your default desktop environment, no configuration changes are necessary. Simply start your X session or log in through *GNOME* display manager like you normally would. If your default desktop environment is *KDE* or a window manager of some sort, check your home directory for the files .Xclients, .xsession, or .xinitrc. If any one of them exists, open it to see if it contains a line similar to **exec startkde**. If it does, remove the executable bit on that file (with **chmod -x .xinitrc/.Xclients/.xsession/**; choose the name of the file that exists in your home directory), and then restart your X session. If it continues to start something other than *GNOME*, create the file .xsession in your home directory (with **touch .xsession**), then edit it with your favorite text editor to read **exec gnome-session**. Save this file, set the executable bit on it (with **chmod +x .xsession**), and then restart your X session.

If that's a bit too brief for you, or if you run into problems, you can find an excellent online guide at http://www.karubik.de/gig/index14.html. If you're not online, there's also a version of it on our cover CD – see the disc interface for details.

And there's more...

You also get a selection of accompanying software that was designed for *GNOME*. It's hilariously called *GNOME + Fifth Toe* and features acres of apps that begin with a letter 'G' (and some that don't).

Here's what you get.
Abi –0.7.14
Atomix –0.4.3
Bombermaze –0.6.5
Dia –0.86
Eog –0.6
Firestarter –0.61
Fpm –0.53
Gabber –0.82
Gal –05
Galleon –0.10
Gedit –0.96
Gfax –0.42
Gimp –1.2.1
Glimmer –1.0.1
Gnome-db –0.2.3
Gnomeicu –0.96
Gnomoko –1.1
Gnorpm
Gnucash
Gnumeric
Gob –1.0.7
Googilzer –0.1

And lots more. Again, instructions to install these can be found on the installation guide on the CD. See disc interface for details.

Other software

As if the Forté, Java, GNOME triumvirate wasn't enough, there's a stack of other goodies on the CD too. Here's a brief look at some of them.

Kernel 2.4.3

As ever, we bring you the latest kernel release on the CD, so here's a brief re-cap on how to update your system to use it.

Installing the kernel

Put the kernel tarball from the directory on the CD into a directory where you have permissions (eg. your home directory) and unpack it:

```
gzip -cd linux-2.4.3.tar.gz | tar xvf -
```

Make sure you have no stale .o files and dependencies lying around:

```
cd linux
make mrproper
```

You should now have the sources correctly installed.

Configuring

Do a **make config** to configure the basic kernel. This needs bash to work and it will search for it in $BASH, /bin/bash and /bin/sh (in that order), so one of those must be correct for it to work. If you want to use a X-based configuration tool (I found this much easier), do a **make xconfig**.

Once you've configured the kernel to your system, do a **make dep** to set up all the dependencies correctly.

Compiling

Make sure you have gcc-2.91.66 (egcs-1.1.2) available. gcc 2.95.2 may also work but is not as safe, and gcc 2.7.2.3 is no longer supported. Also remember to upgrade your binutils package (for as/ld/nm and company) if necessary.

Do a **make bzImage** to create a compressed kernel image (remember Linux is case sensitive). If you want to make a boot disk (without root filesystem or LILO), insert a floppy in your A: drive, and do a **make bzdisk**. It is also possible to do **make install** if you have lilo installed to suit the kernel makefiles, but you may want to check your particular lilo setup first.

To do the actual install you have to be logged in as root, but none of the normal builds should require that. If you configured any of the parts of the kernel as modules, you will have to do **make modules** followed by **make modules_install**.

If you need more help, read Documentation/modules.txt, which has a comprehensive explanation of how to use the modules.

In order to boot your new kernel, you'll need to copy the kernel image (found in .../linux/arch/i386/boot/bzImage after compilation) to the place where your regular bootable kernel is found.

If you boot Linux from the hard drive, chances are you'll be using LILO which uses the kernel image as specified in the file /etc/lilo.conf. The kernel image file is usually /vmlinuz, /boot/vmlinuz, /bzImage or /boot/bzImage. To use the new kernel, save a copy of the old image and copy the new image in its place. Then, you must rerun LILO to update the loading map! If you don't, you won't be able to boot the new kernel image.

Reinstalling LILO is usually a matter of running /sbin/lilo. You may wish to edit /etc/lilo.conf to specify an entry for your old kernel image (say, /vmlinux.old) in case the new one doesn't work out the way it should. Have a scroll through the LILO docs for more information.

After reinstalling LILO, you should be all set. Shutdown the system, reboot, and enjoy!

CompuPic

CompuPic is a free-for-personal use image cataloguing application that is great for use in conjunction with image manipulators such as The GIMP.

As well as cataloguing, CompuPic provides facilities for viewing and making small adjustments to your images and also sharing them with other users on the 'net. See http://www.photodex.com for more.

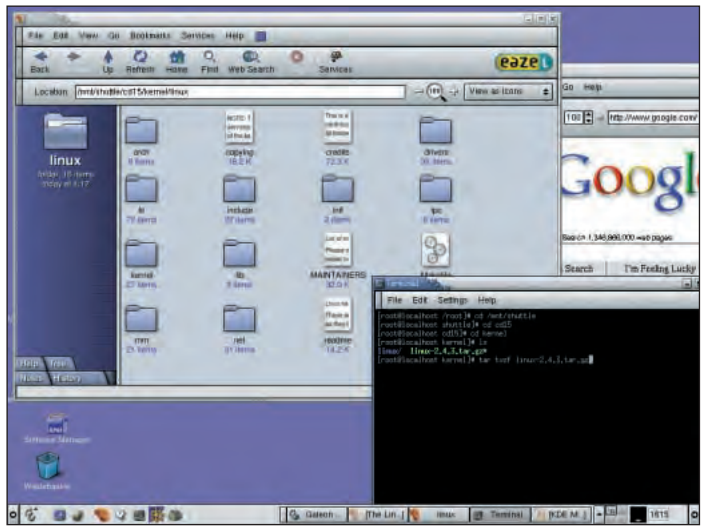

Updating your kernel should improve your PC's performance.

There's a new release of Mozilla, and it's on the CD!

KDE 2.1.1

Yes, no sooner then you'd got *KDE 2.1* installed on your system, along comes *2.1.1*. The major changes in this version seem to be in the *Konqueror* web browser, so it is a worthwhile upgrade. In the time between sending the CD to press and writing this, *KDE* have updated kdelibs to 2.1.2, which means a big bugfix of the core KDE libraries. Look out for this essential update and others on your next *Linux Format*.

Mozilla 0.8.1

The *Mozilla* project lumbers on towards its goal, and it's spawned another release along the way. Mozilla 0.8.1 features a new, improved Chatzilla IRC client and a similarly overhauled JavaScript console.

There's also improved theme switching and removal, a new hierarchical history – allowing you to sort history by Day, Site and Page using an expanding list similar to that found in Internet Explorer. Other changes include improved mail message filtering, better FTP performance, mail message keyboard navigation improvements and a shrinkable progress window.

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

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

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

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

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

Editorial Contributors
Charlie Stross, Mike Saunders, David Coulson, Jono Bacon, Jon Kent, Paul Cavanagh, John Walker, Biagio Lucini, Chris Howells, Andrew Arensburger

ART CONTRIBUTORS
Photography Rick Buettner, Louise Parker, Gavin Roberts
Illustration Paul Bateman, Chris Winn

ADVERTISING
Senior Advertising Manager Brian Ainge
Deputy Ad 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, S Hobbs, M Stapleton, C Stenner
Supervisor Sarah Orchard

MARKETING AND PROMOTIONS
Marketing Manager Gill Stevenson
Marketing Exec Liz Britton
Subscriptions Manager Julie Sewell

PRODUCTION
Production Manager Lou Reffell
Production Co-ordinator Ian Wardle
Ad Designer David Mathews
Group Production Manager Judith Green

MANAGEMENT
Publisher Sheena Pittaway
Group Publisher John Weir
Circulation Director Sue Hartley
Managing Director Mike Frey

DISTRIBUTION AND CIRCULATION
Circulation Manager Richard Jefferies
richard.jefferies@futurenet.co.uk

Distributed by Seymour Distribution, 86 Newman Street, London W1P 3LD
Tel +44 (0)207 9076000

Overseas Licences
Tel +44(0)1225 442244 **Fax** +44 (0)1225 732384

Contact Details
Linux Format, 30 Monmouth Street, Bath BA1 2BW
Tel +44 (0)1225 442244 **Email** linuxformat@futurenet.co.uk

Subscriptions and Mail Order
Phone +44 (0)1458 271178. See page 88
Email subs@futurenet.co.uk

Next month

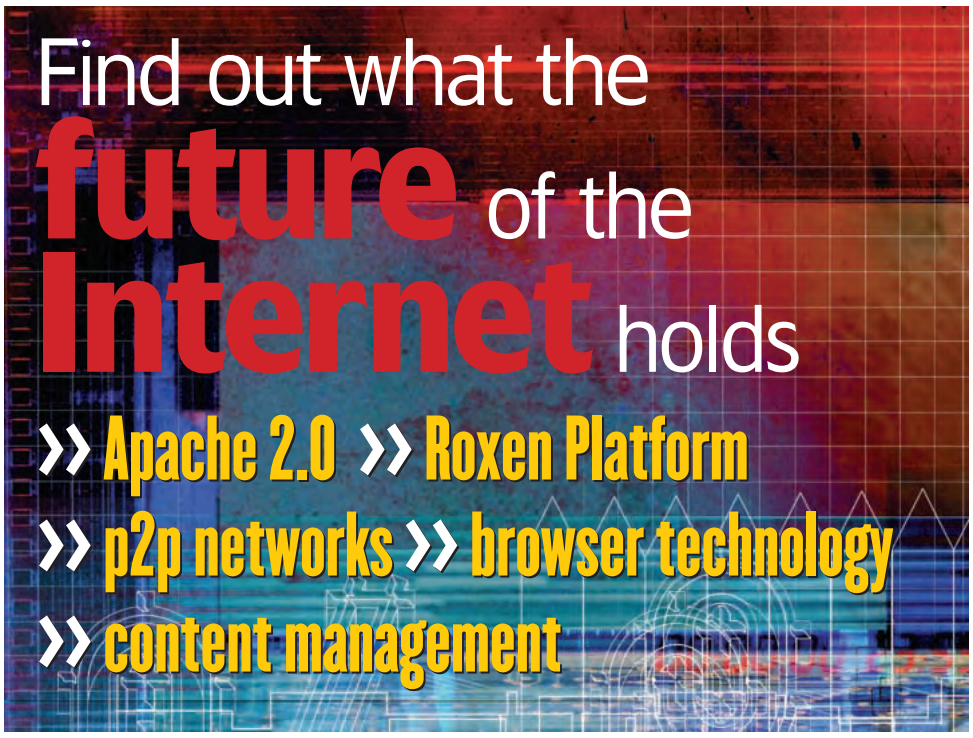

» Latest Distros on test
Red Hat 7.1, SuSE 7.1 PPC and Mandrake 8.0 reviewed and rated

» Game on!
The mysteries of OpenGL games programming revealed in a new series

» OS X
Does Apple's new slant on Unix have anything to teach us?

» Latest OpenOffice, Mozilla, Opera and TurboPrint on the CD, plus latest kernel source, HotPicks and More

LINUXformat
Issue 16 on sale Wednesday 20 June

The UK's premier forum for Linux computing

Linux EXPO 2001

EXHIBITION & CONFERENCE

Register FREE at www.linuxexpo.co.uk

London • Olympia • UK • 4th - 5th July 2001

Now in its third year, Linux Expo, taking place at Olympia, London is firmly established as the UK's largest Exhibition and Conference dedicated to this open source revolution. It brings together all the major players in the Linux computing sphere to provide the definitive review of all the latest products and services emerging around this rapidly developing technology.

If you are evaluating the importance and impact of Linux technology on your IT infrastructure, you can meet over 100 international players who are leading the Linux revolution. From development tools, e-commerce, security and Internet tools to Linux based training and support; this is your unique opportunity to discuss your Linux issues directly with the industry leaders.

Visit **www.linuxexpo.co.uk** to register for your free admission badge.

Show sponsored by

IBM CALDERA SuSE

Online sponsor

ZDNet UK www.zdnet.co.uk

PLUS, RUNNING ALONGSIDE

THE FIRST LINUX CONFERENCE FOR IT PROFESSIONALS IN THE LARGE ENTERPRISE.

ENTERPRISE LINUX INSTITUTE

July 3-5, 2001
Olympia Conference Centre.

Whether you're currently using or planning to integrate/migrate to Linux, this conference is designed for serious enterprise Linux business professionals, the Enterprise Linux Institute Conference featured tracks include: Linux Fundamentals, Linux for e-business and Enterprise strength Linux.

For additional information and to register: **www.elxi.co.uk**
or email: **enquiries@101com.co.uk**

IT EVENTS

a 101 communications Company

The only established Linux Computing exhibition in the UK